



(12) **United States Patent**
Tomlin et al.

(10) **Patent No.:** **US 9,268,682 B2**
(45) **Date of Patent:** **Feb. 23, 2016**

(54) **METHODS, DEVICES AND SYSTEMS FOR
PHYSICAL-TO-LOGICAL MAPPING IN
SOLID STATE DRIVES**

(71) Applicants: **Western Digital Technologies, Inc.**,
Irvine, CA (US); **Skyera, LLC**, San
Jose, CA (US)

(72) Inventors: **Andrew J. Tomlin**, San Jose, CA (US);
Rodney N. Mullendore, San Jose, CA
(US); **Justin Jones**, Redwood City, CA
(US); **Radoslav Danilak**, Cupertino, CA
(US)

(73) Assignees: **Skyera, LLC**, San Jose, CA (US);
Western Digital Technologies, Inc.,
Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 505 days.

(21) Appl. No.: **13/645,822**

(22) Filed: **Oct. 5, 2012**

(65) **Prior Publication Data**
US 2014/0101369 A1 Apr. 10, 2014

(51) **Int. Cl.**
G06F 12/02 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 12/0246** (2013.01); **G06F 2212/7201**
(2013.01)

(58) **Field of Classification Search**
CPC G06F 9/466; G06F 9/467; G06F 11/1474;
G06F 17/30345; G06F 12/0246; G06F
2212/7201

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,363,421 B2	4/2008	Di Sena et al.
7,412,585 B2	8/2008	Uemura
7,441,067 B2	10/2008	Gorobets et al.
7,509,471 B2	3/2009	Gorobets
7,610,438 B2	10/2009	Lee et al.
8,095,851 B2	1/2012	Diggs et al.
8,407,449 B1	3/2013	Colon et al.
8,489,854 B1	7/2013	Colon et al.
2004/0088474 A1	5/2004	Lin
2004/0109376 A1	6/2004	Lin
2006/0155917 A1	7/2006	Di Sena et al.
2007/0016721 A1	1/2007	Gay

(Continued)

OTHER PUBLICATIONS

Written Opinion & International Search Report dated: Jan. 16, 2014
from Patent Application Serial No. PCT/US2013/062723 11 pages.

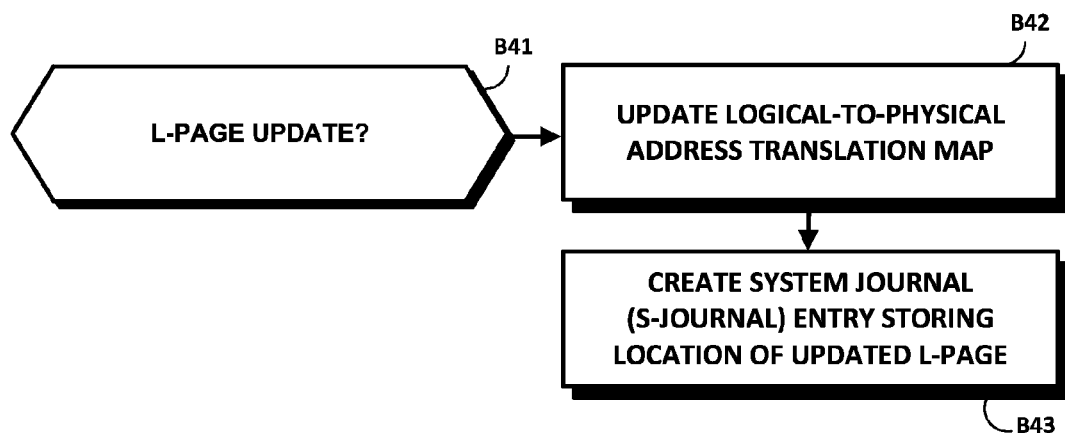
Primary Examiner — Gary W Cygiel

(74) *Attorney, Agent, or Firm* — Wilmer Cutler Pickering
Hale and Dorr LLP

(57) **ABSTRACT**

A data storage device comprises a plurality of non-volatile memory devices storing physical pages, each stored at a pre-determined physical location. A controller may be coupled to the memory devices and configured to access data stored in a plurality of logical pages (L-Pages), each associated with an L-Page number that enables the controller to logically reference data stored in the physical pages. A volatile memory may comprise a logical-to-physical address translation map that enables the controller to determine a physical location, within the physical pages, of data stored in each L-Page. The controller may be configured to maintain, in the memory devices, journals defining physical-to-logical correspondences, each journal covering a predetermined range of physical pages and comprising a plurality of entries that associate one or more physical pages to each L-Page. The controller may read the journals upon startup and rebuild the address translation map from the read journals.

45 Claims, 17 Drawing Sheets



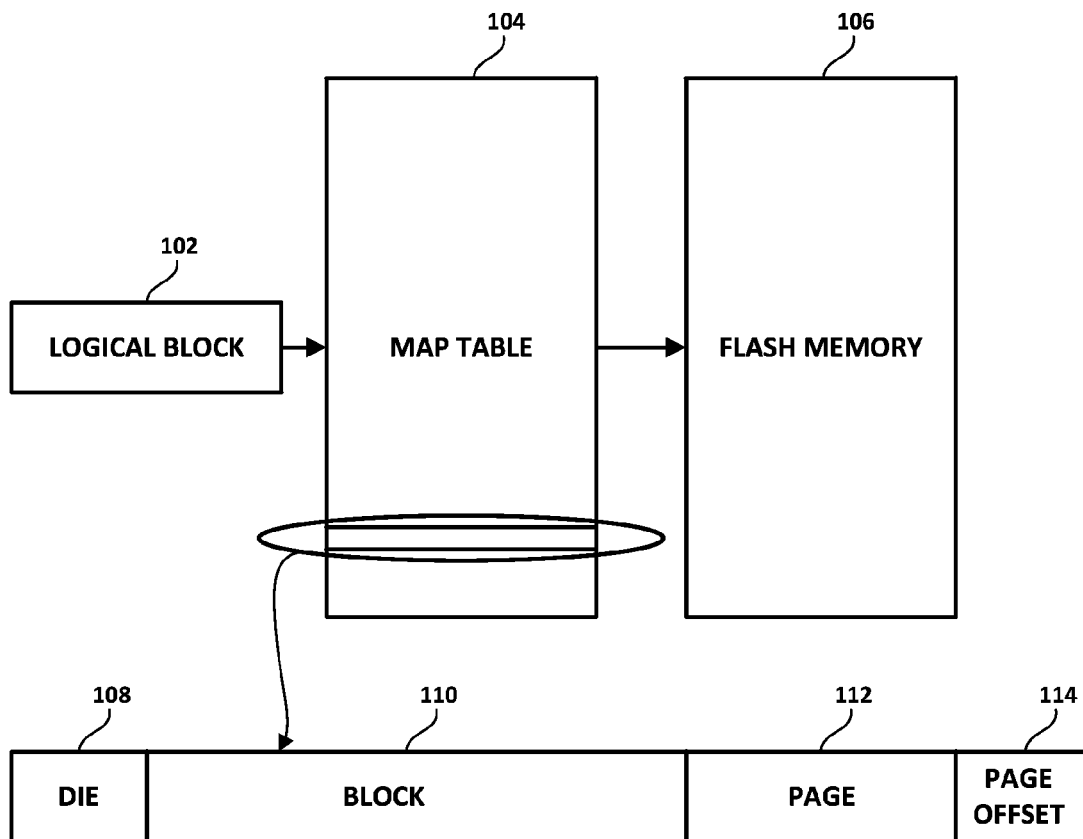
(56)

References Cited

U.S. PATENT DOCUMENTS

2007/0094445	A1	4/2007	Trika et al.	2010/0211851	A1	8/2010	Dixon
2007/0204128	A1	8/2007	Lee et al.	2010/0241928	A1	9/2010	Kim et al.
2007/0300037	A1	12/2007	Rogers et al.	2010/0268871	A1	10/2010	Lee et al.
2008/0104308	A1	5/2008	Mo et al.	2010/0306451	A1	12/2010	Johnson
2008/0256287	A1	10/2008	Lee et al.	2011/0055458	A1	3/2011	Kuehne
2008/0282024	A1	11/2008	Biswas et al.	2011/0066788	A1	3/2011	Eleftheriou et al.
2009/0019218	A1	1/2009	Sinclair et al.	2011/0072194	A1	3/2011	Forhan et al.
2009/0049229	A1	2/2009	Honda et al.	2011/0072333	A1	3/2011	Kuo
2009/0119353	A1	5/2009	Oh et al.	2011/0099321	A1	4/2011	Haines et al.
2009/0150599	A1	6/2009	Bennett	2011/0099350	A1	4/2011	Feldman et al.
2009/0177931	A1	7/2009	Song et al.	2011/0119464	A1	5/2011	Karr et al.
2009/0222643	A1	9/2009	Chu	2011/0191566	A1	8/2011	Takamiya et al.
2009/0240873	A1	9/2009	Yu et al.	2011/0191654	A1	8/2011	Rub
2009/0241008	A1	9/2009	Kim et al.	2011/0225346	A1	9/2011	Goss et al.
2009/0241009	A1	9/2009	Kong et al.	2011/0252289	A1	10/2011	Patapoutian et al.
2009/0259805	A1	10/2009	Kilzer et al.	2011/0283049	A1	11/2011	Kang et al.
2009/0292972	A1	11/2009	Seol et al.	2011/0320915	A1	12/2011	Khan
2010/0030999	A1	2/2010	Hinz	2012/0023387	A1	1/2012	Wang et al.
2010/0174851	A1	7/2010	Leibowitz et al.	2012/0072807	A1	3/2012	Cornwell et al.
				2012/0173795	A1	7/2012	Schuette et al.
				2012/0226887	A1	9/2012	Culley et al.
				2013/0091322	A1*	4/2013	Wang et al. 711/103

* cited by examiner

**FIG. 1***(Prior Art)*

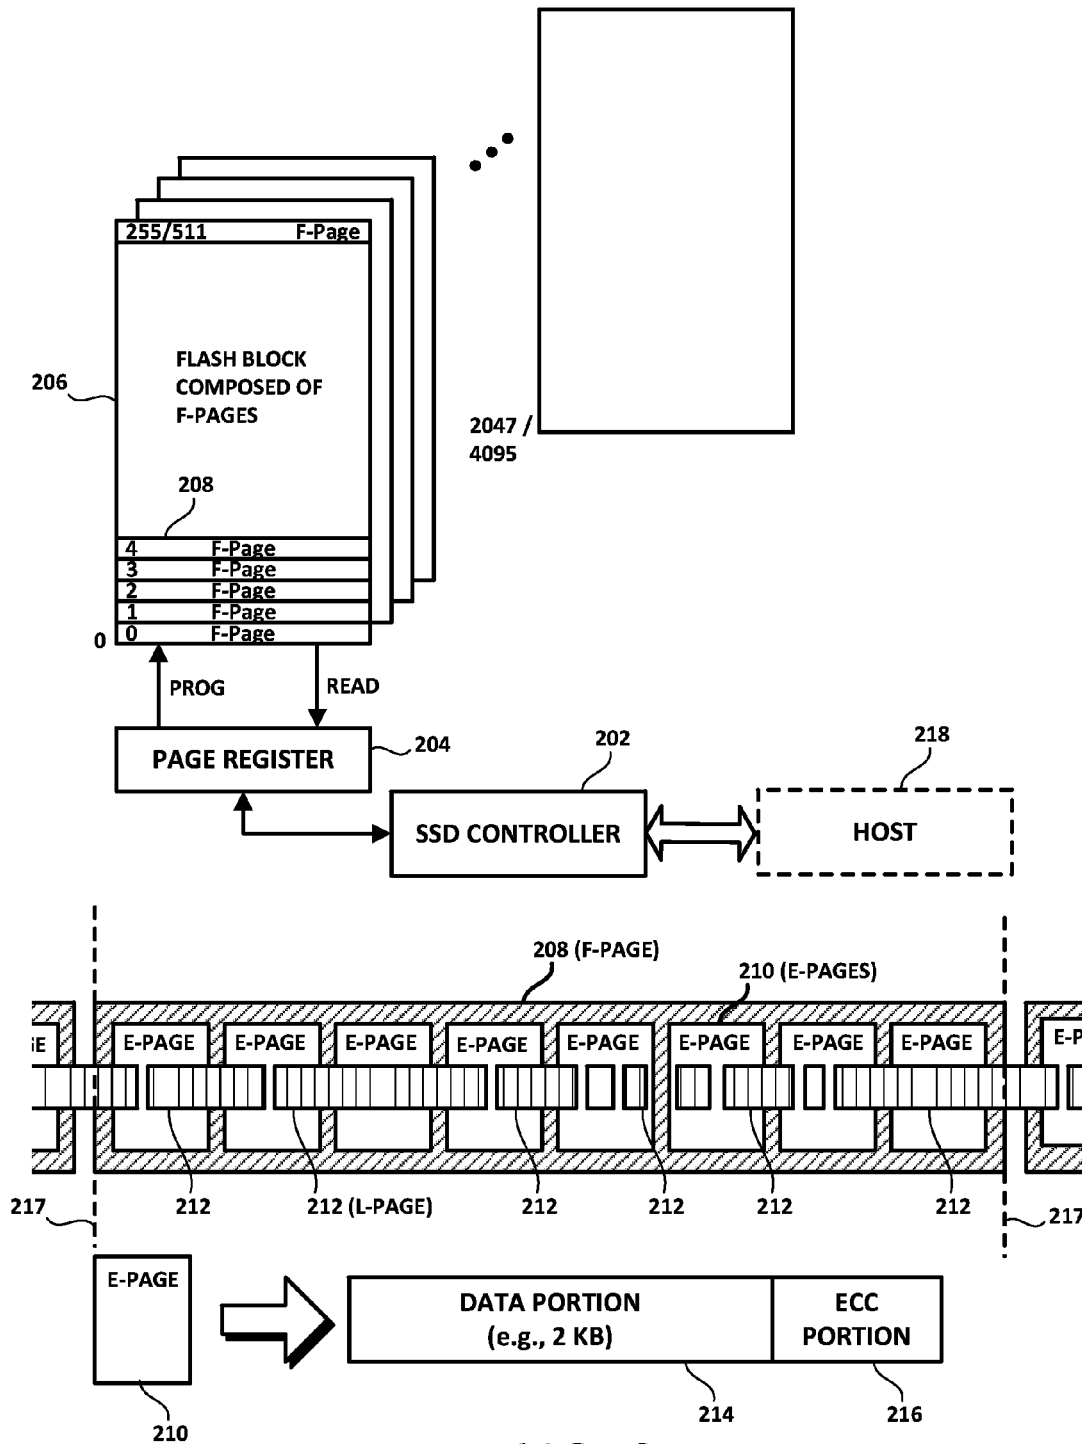
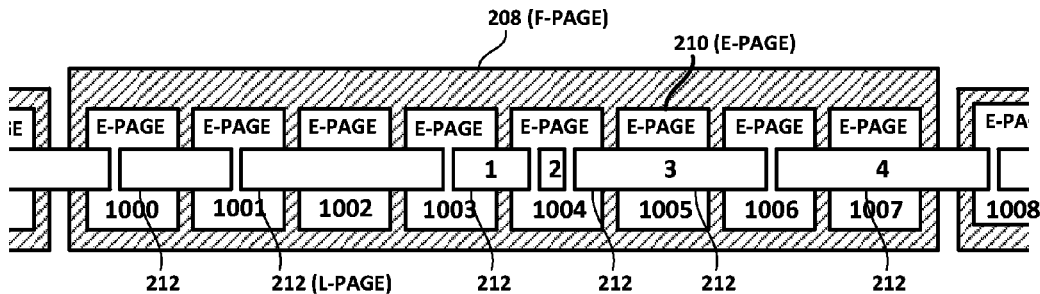


FIG. 2



<u>L-PAGE #</u>				<u>E-PAGES NEEDED</u>
L-Page 1 =	E-Page 1003	Offset 800	Len 1,624	1003,1004
L-Page 2 =	E-Page 1004	Offset 400	Len 696	1004
L-Page 3 =	E-Page 1004	Offset 1,120	Len 4,096	1004, 1005, 1006
L-Page 4 =	E-Page 1006	Offset 1,144	Len 3,128	1006, 1007, 1008

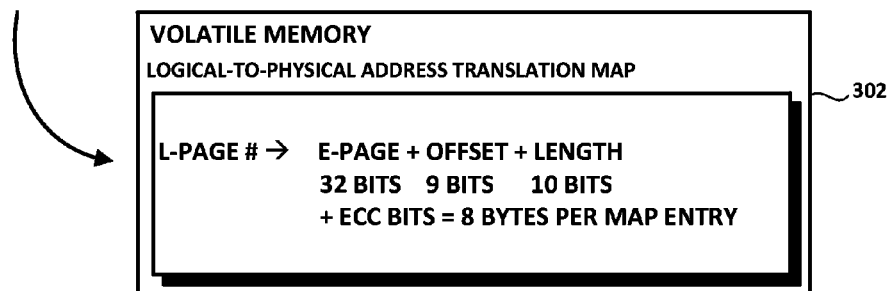


FIG. 3

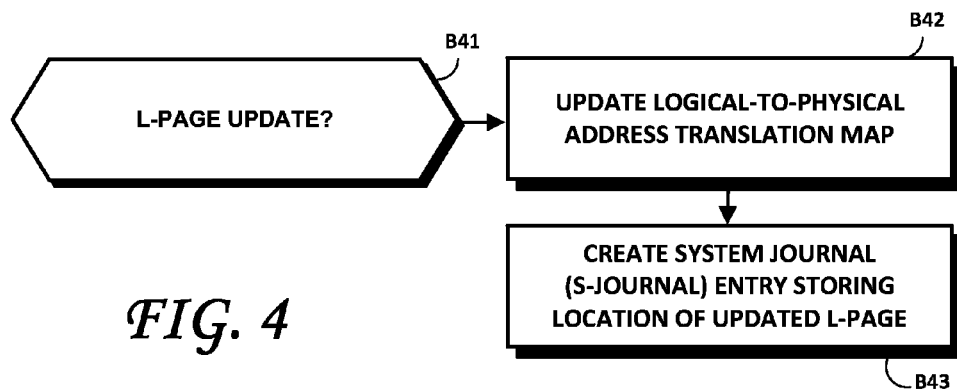
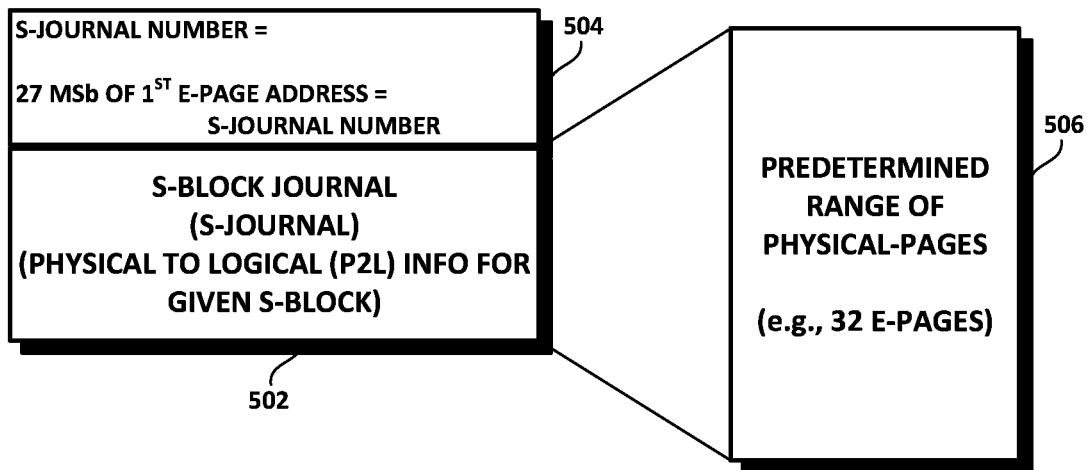
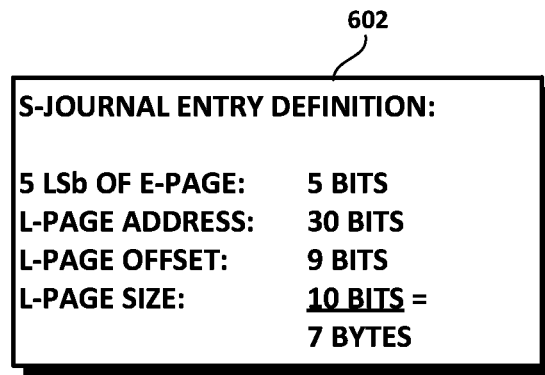


FIG. 4

*FIG. 5**FIG. 6*

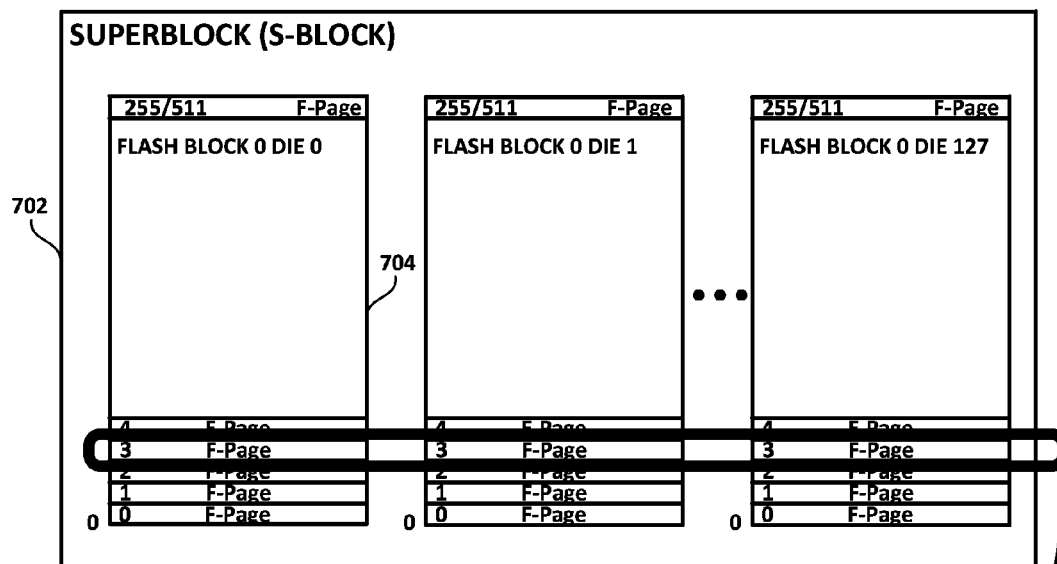


FIG. 7

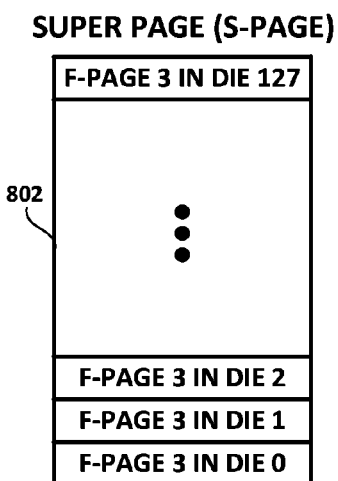
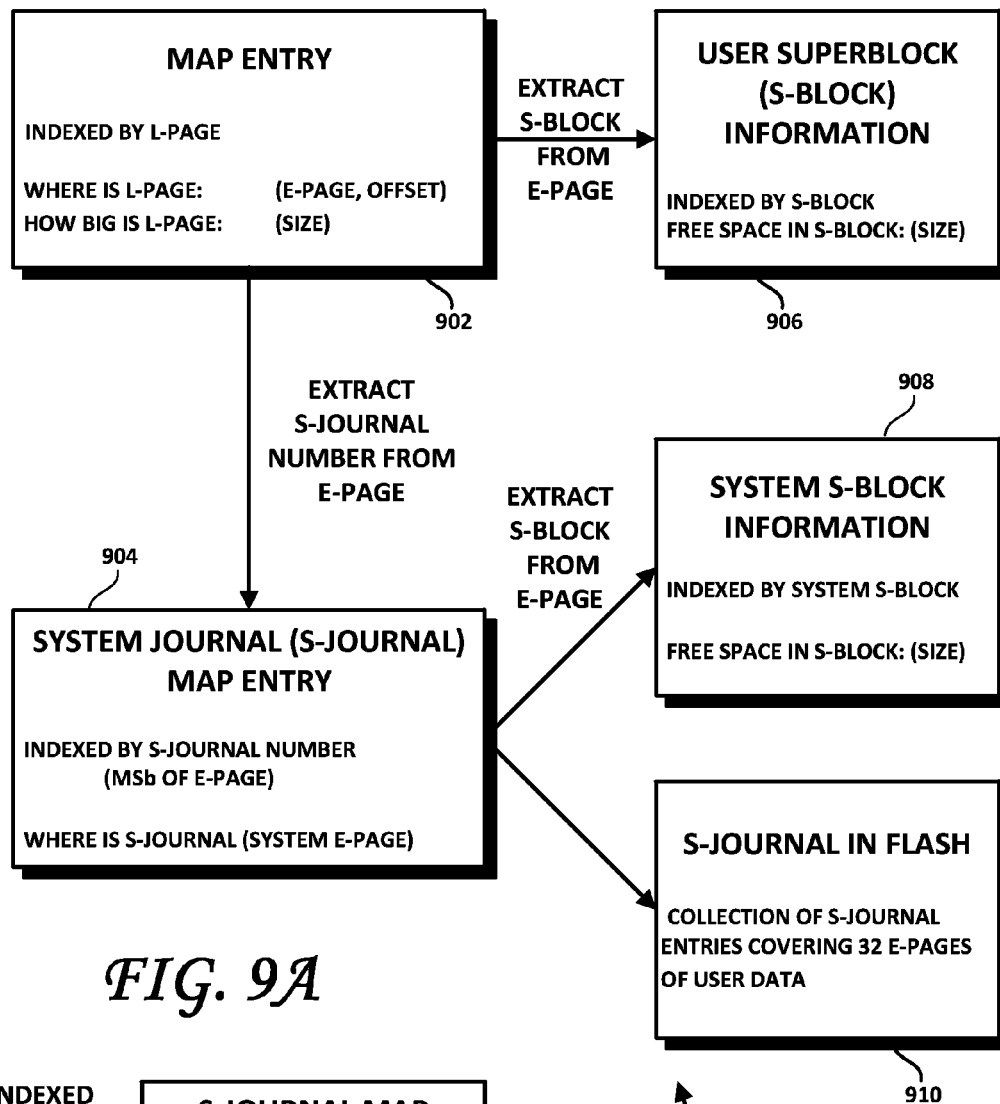
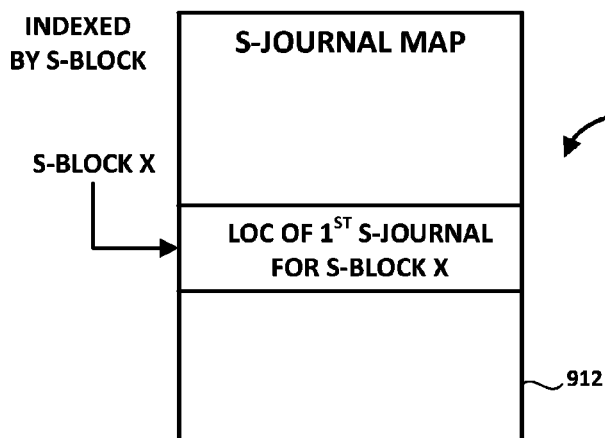
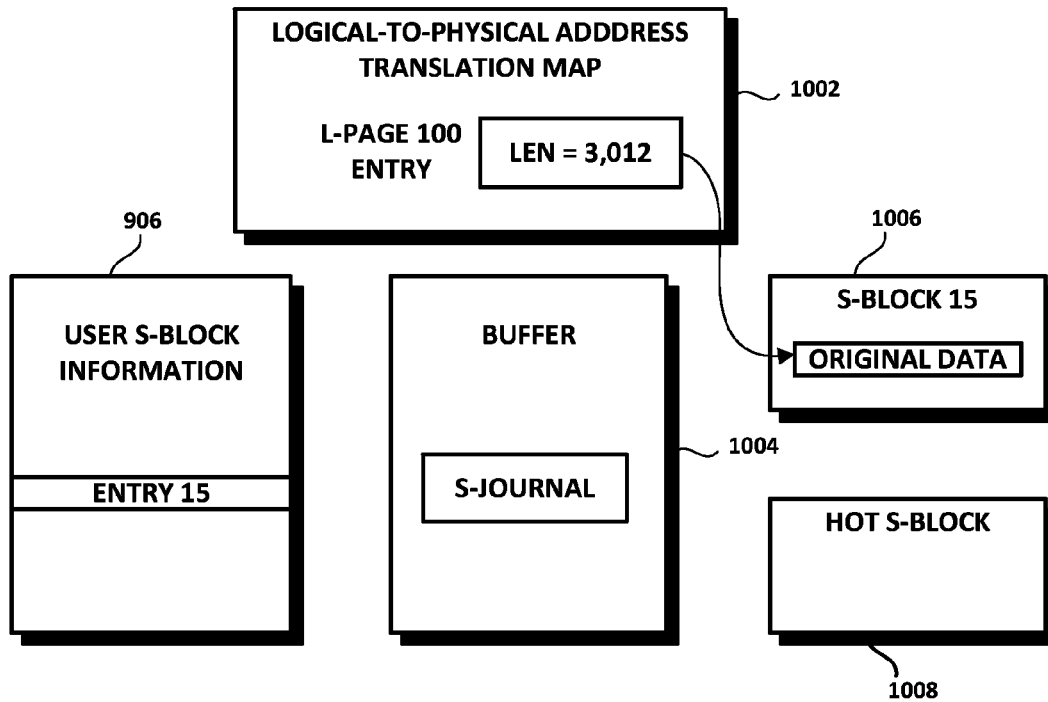
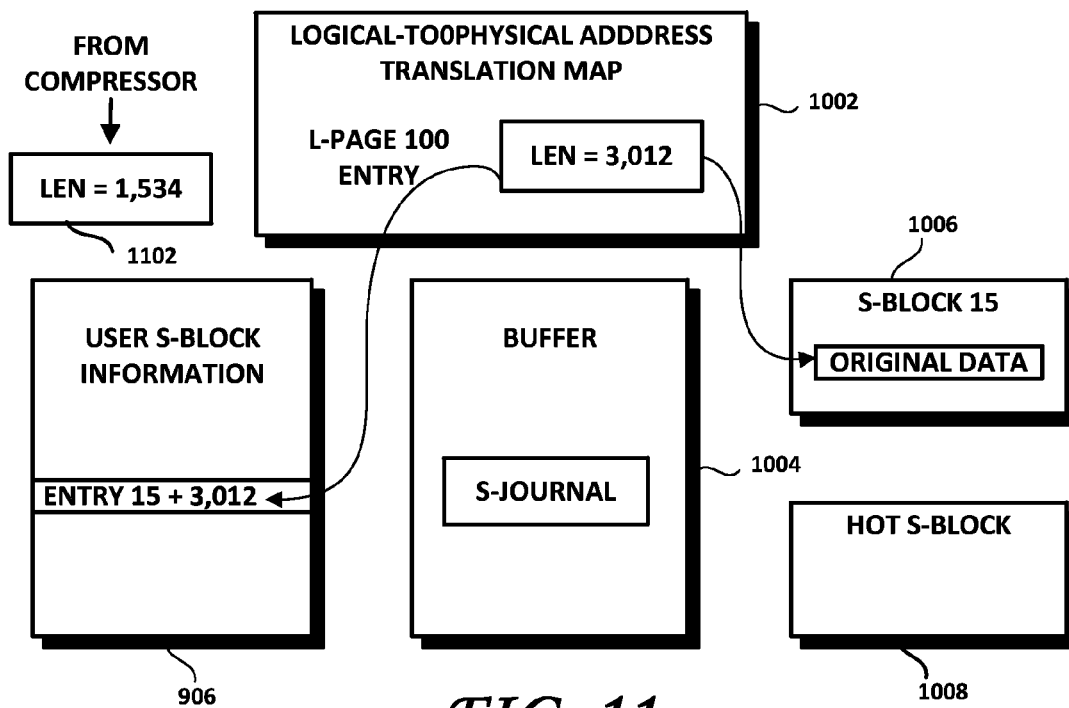
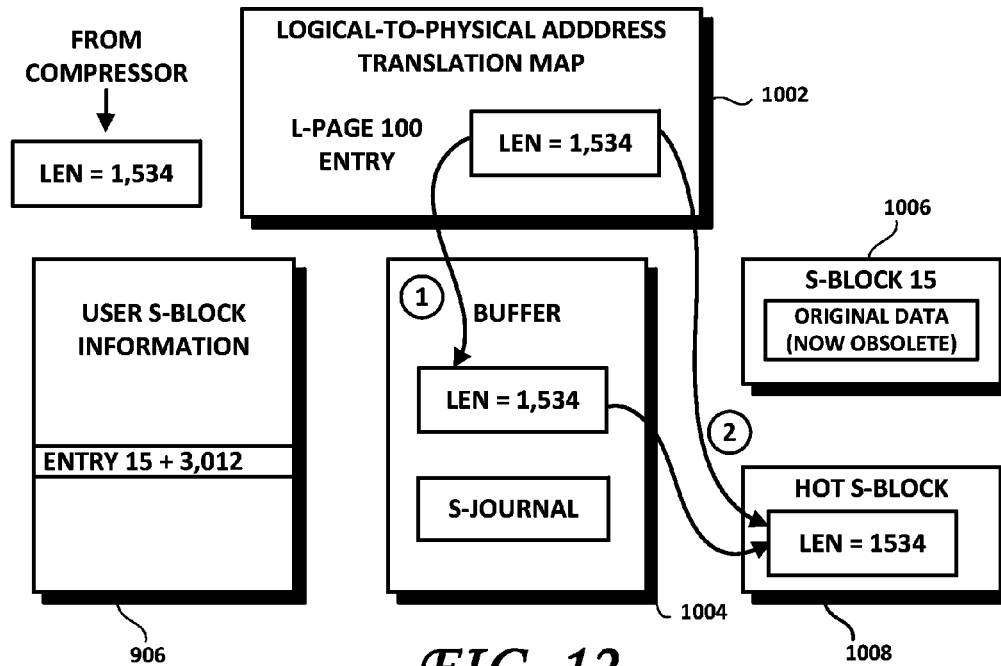
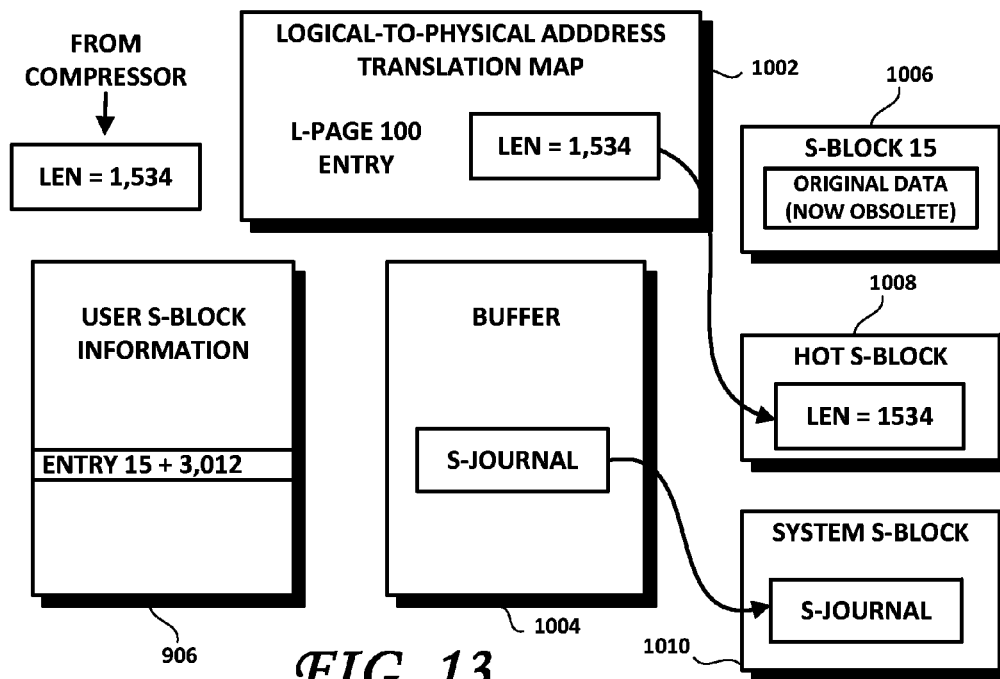


FIG. 8

*FIG. 9A**FIG. 9B*

*FIG. 10**FIG. 11*

*FIG. 12**FIG. 13*

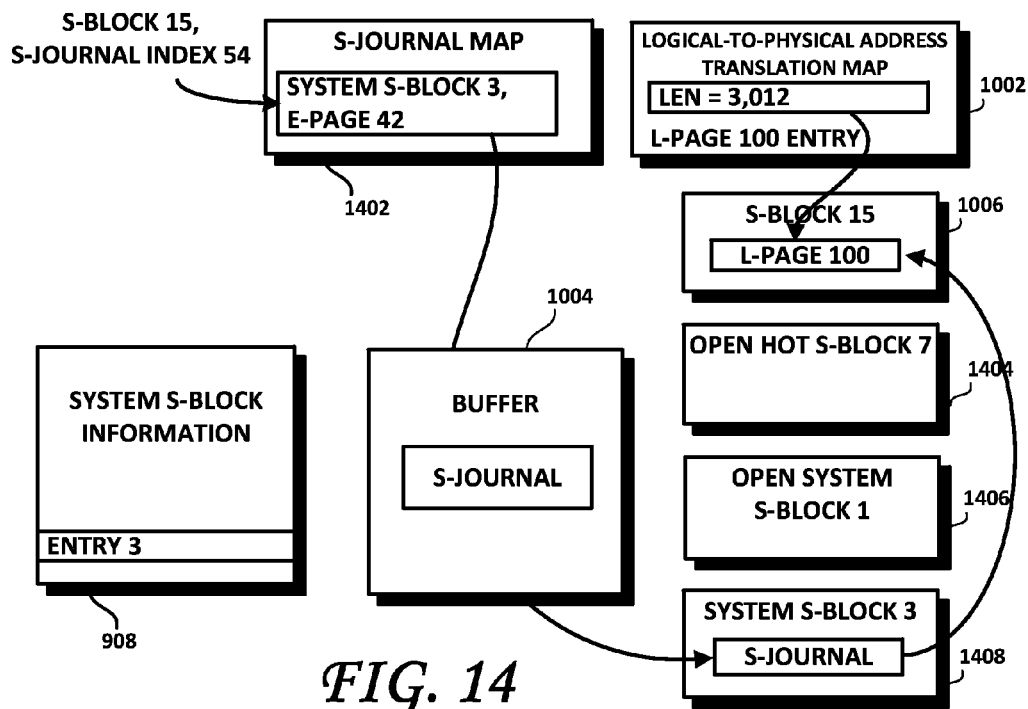


FIG. 14

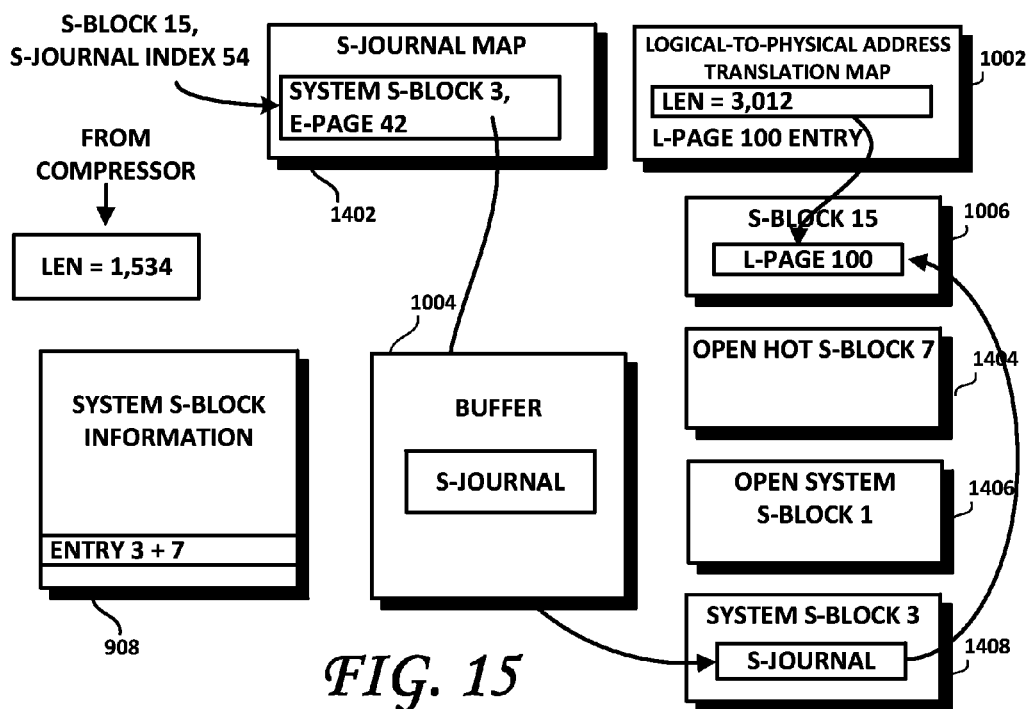
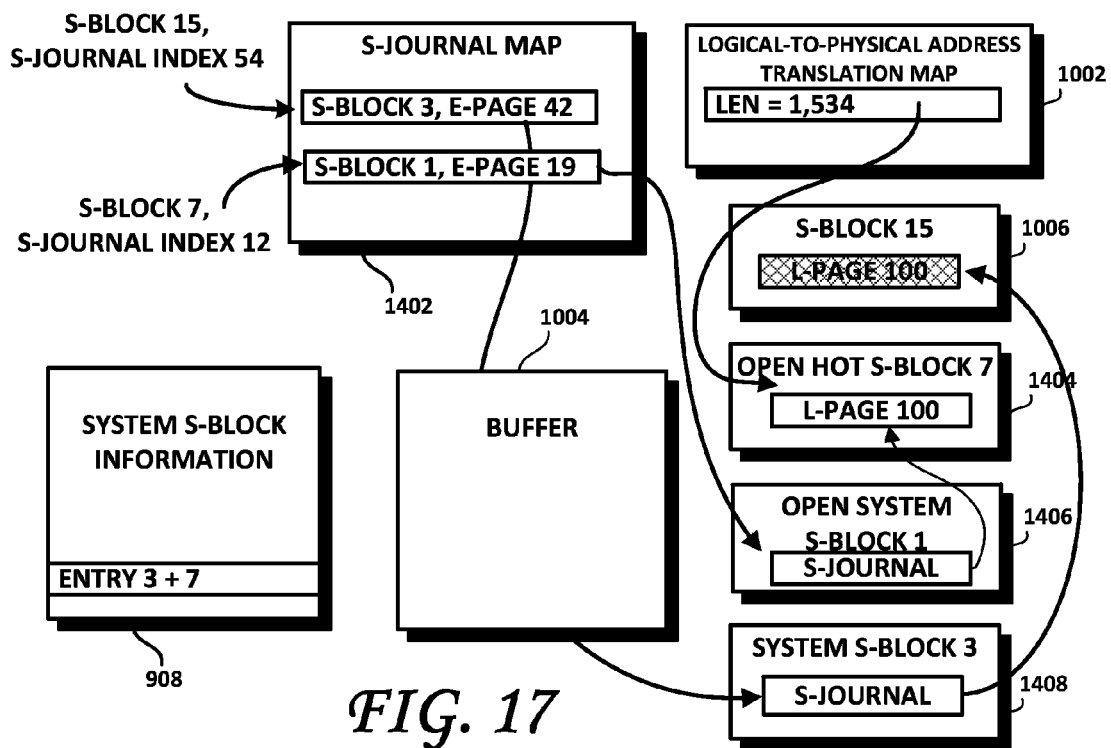
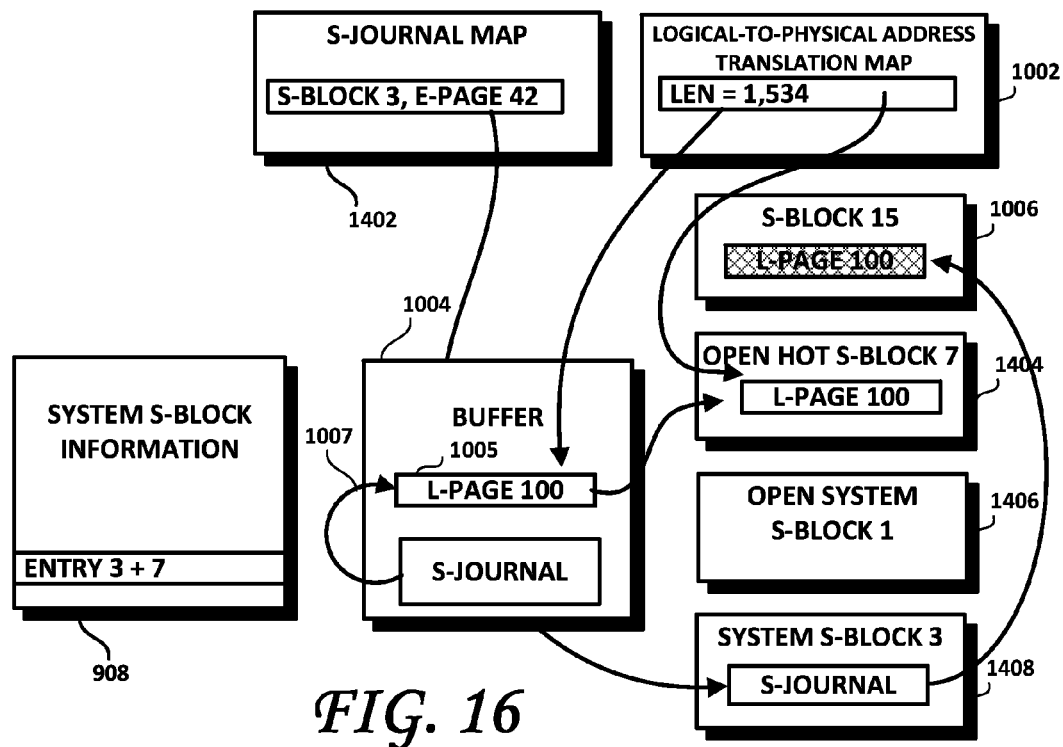
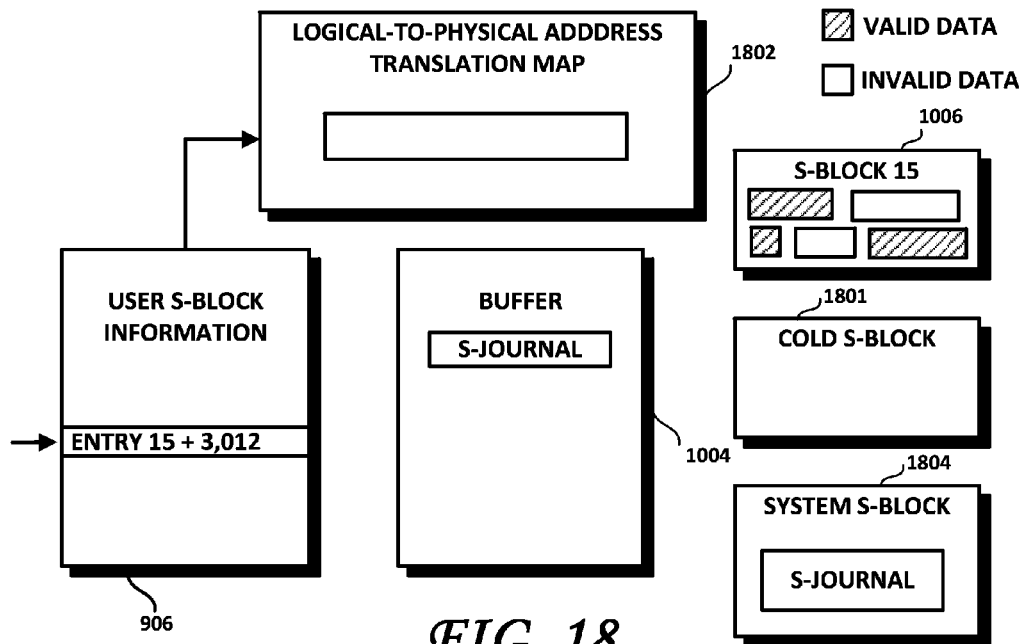
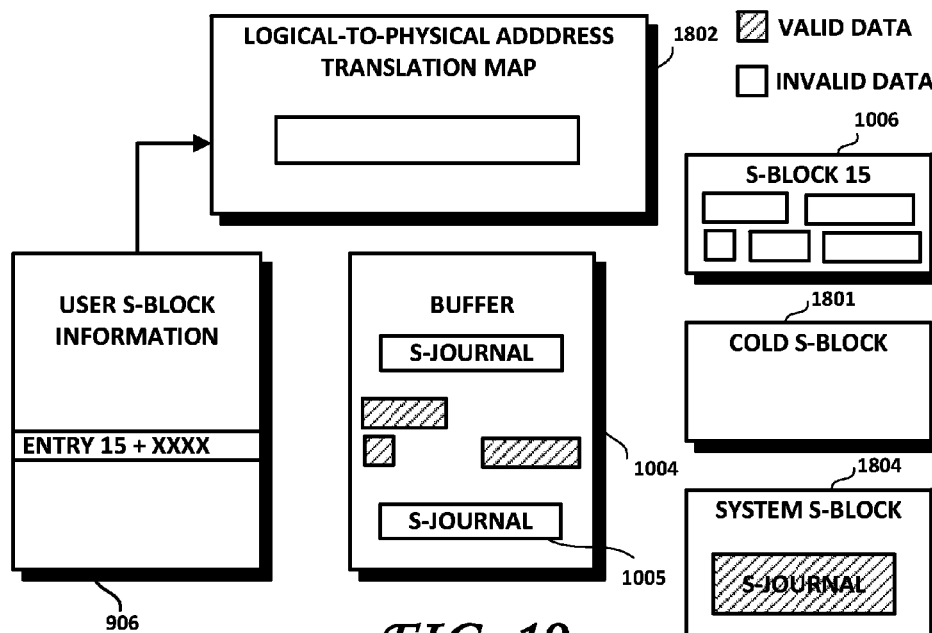
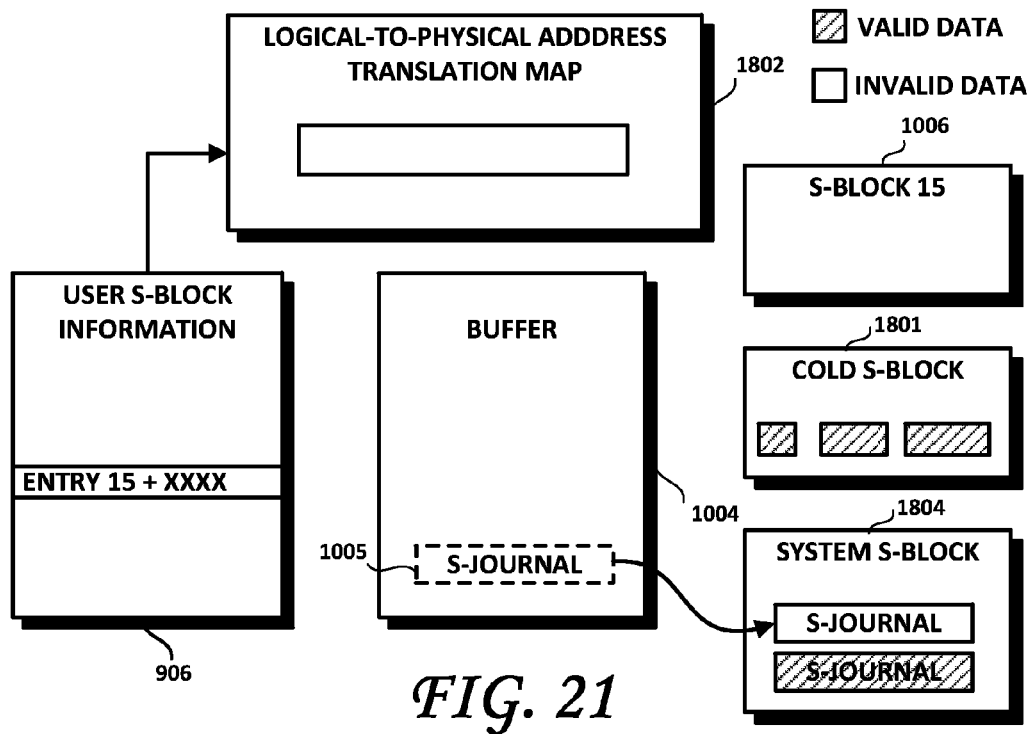
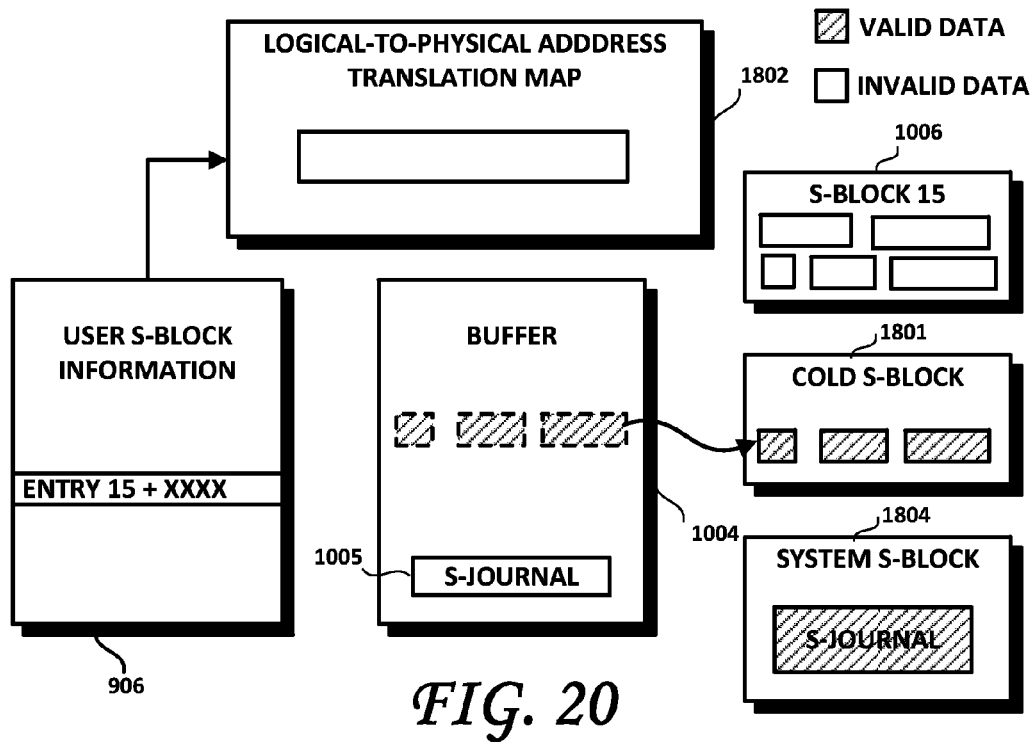


FIG. 15



*FIG. 18**FIG. 19*



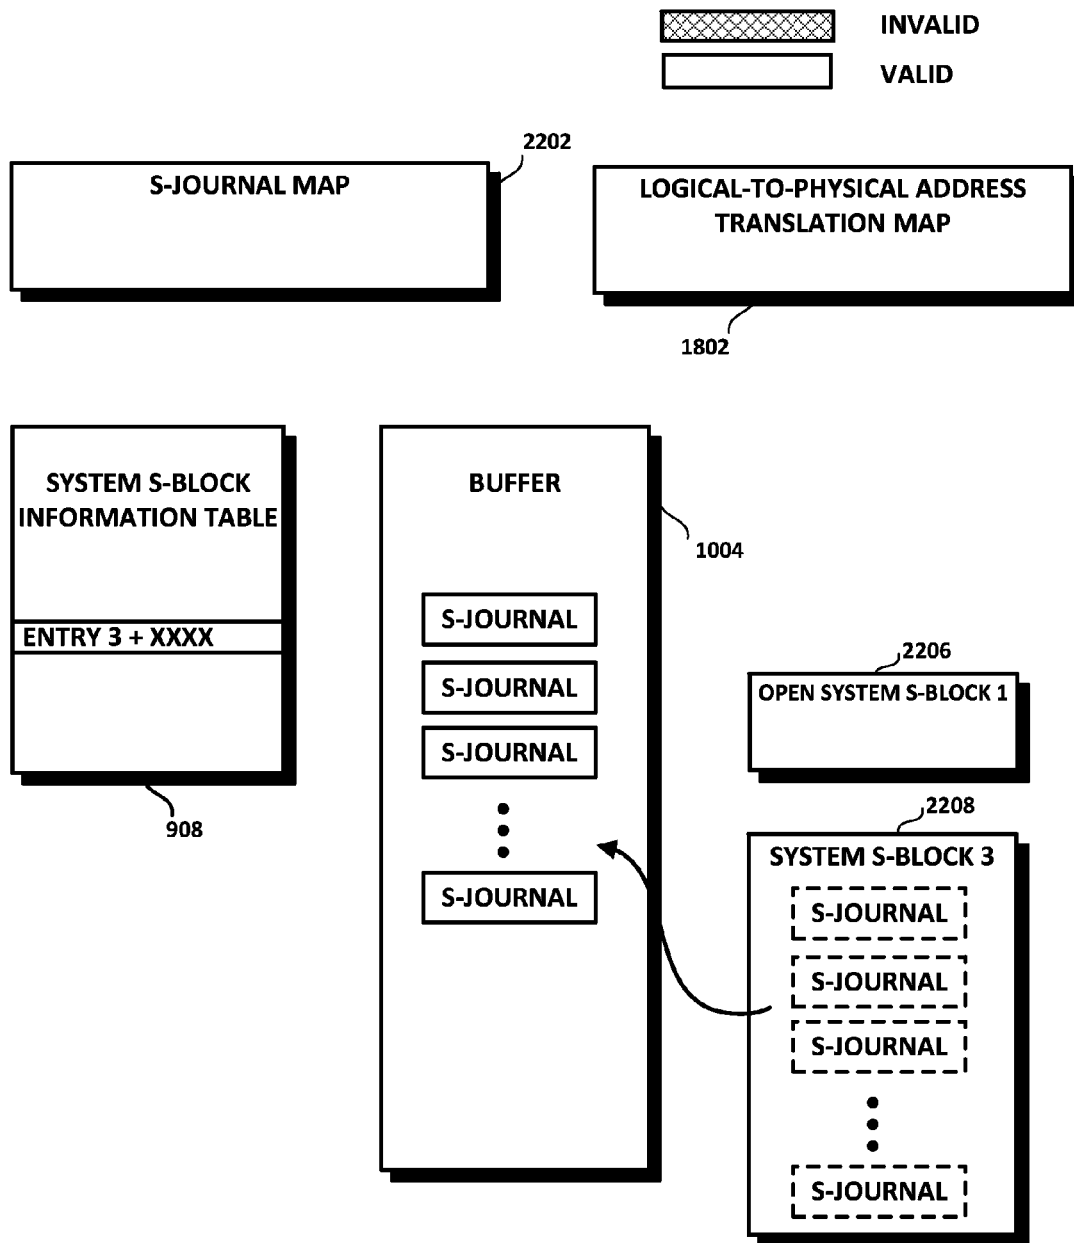


FIG. 22

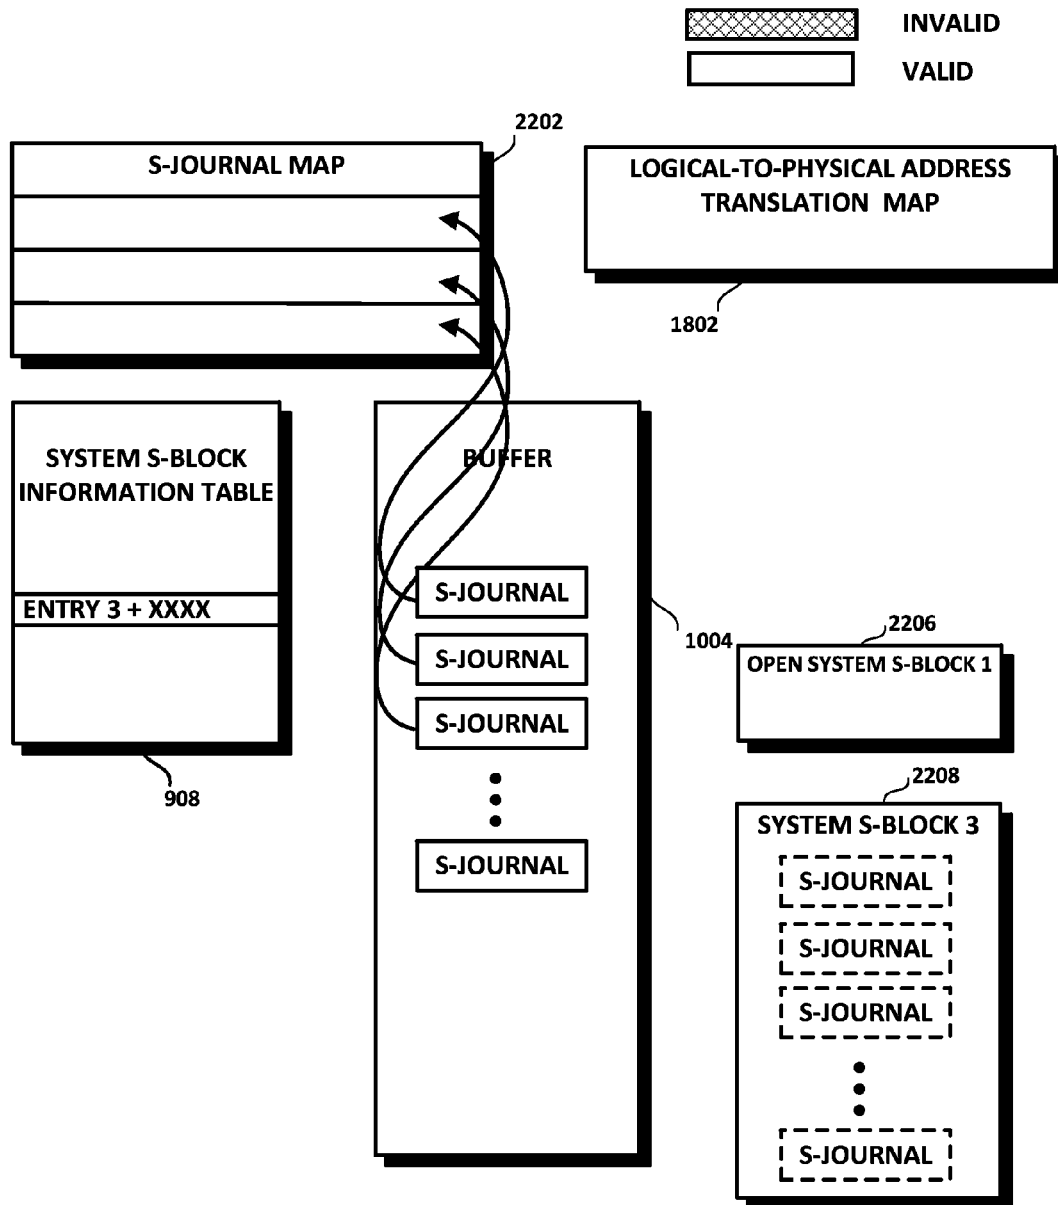


FIG. 23

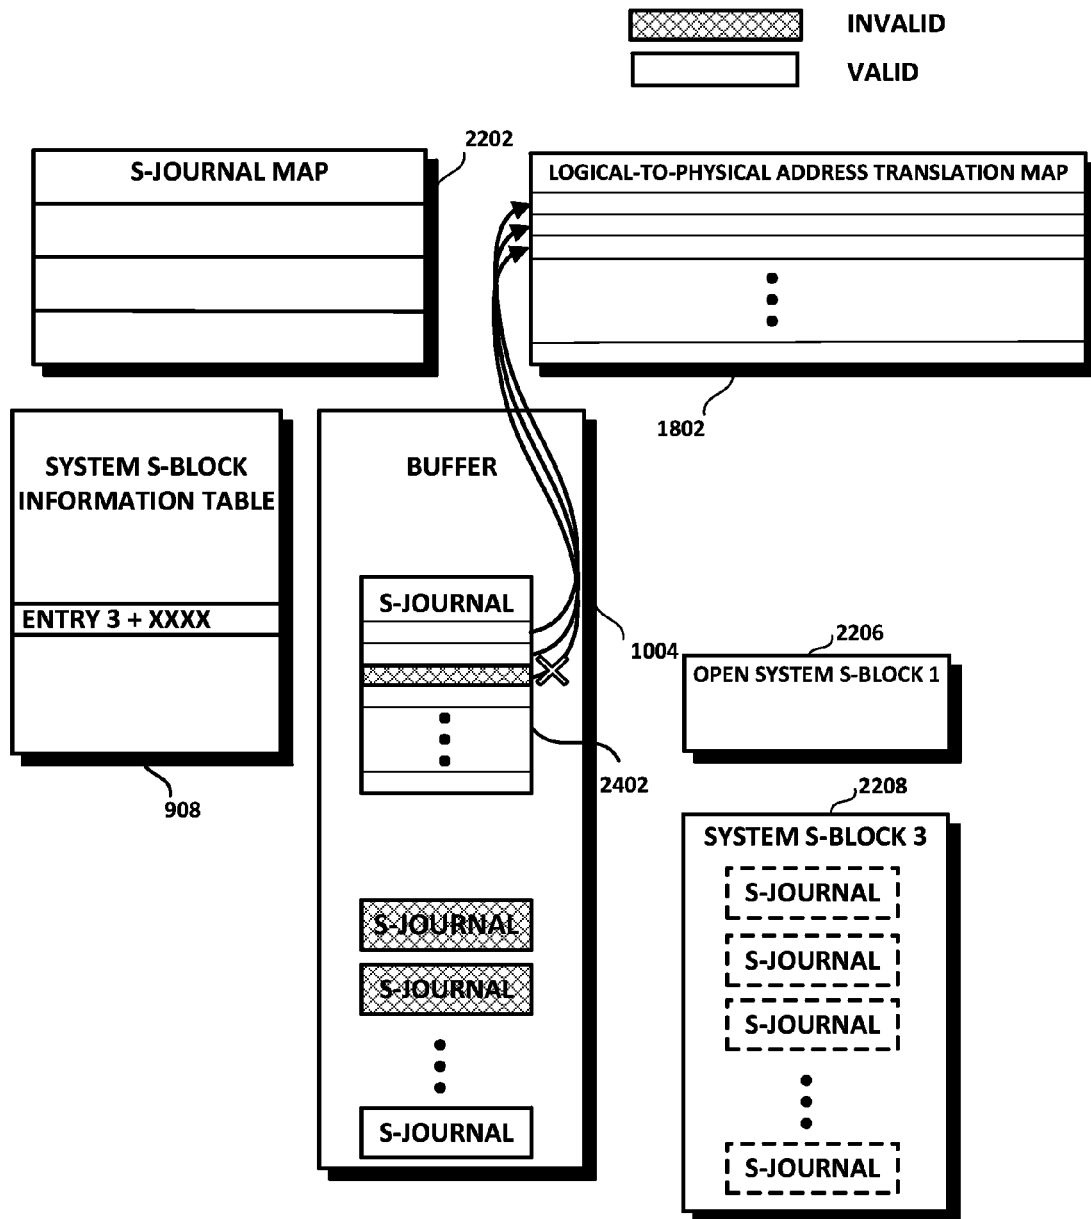


FIG. 24

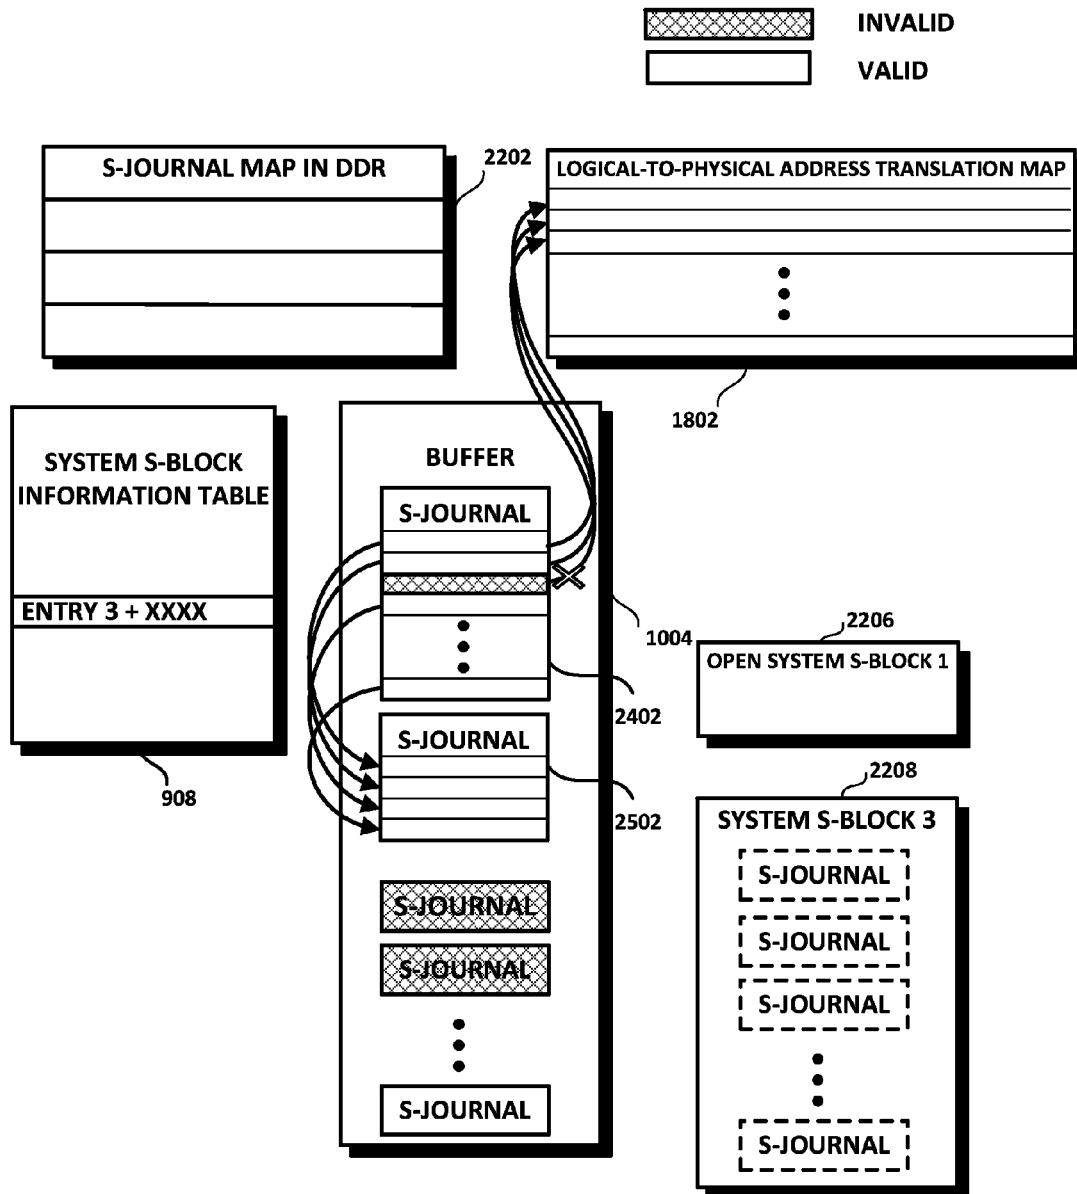


FIG. 25

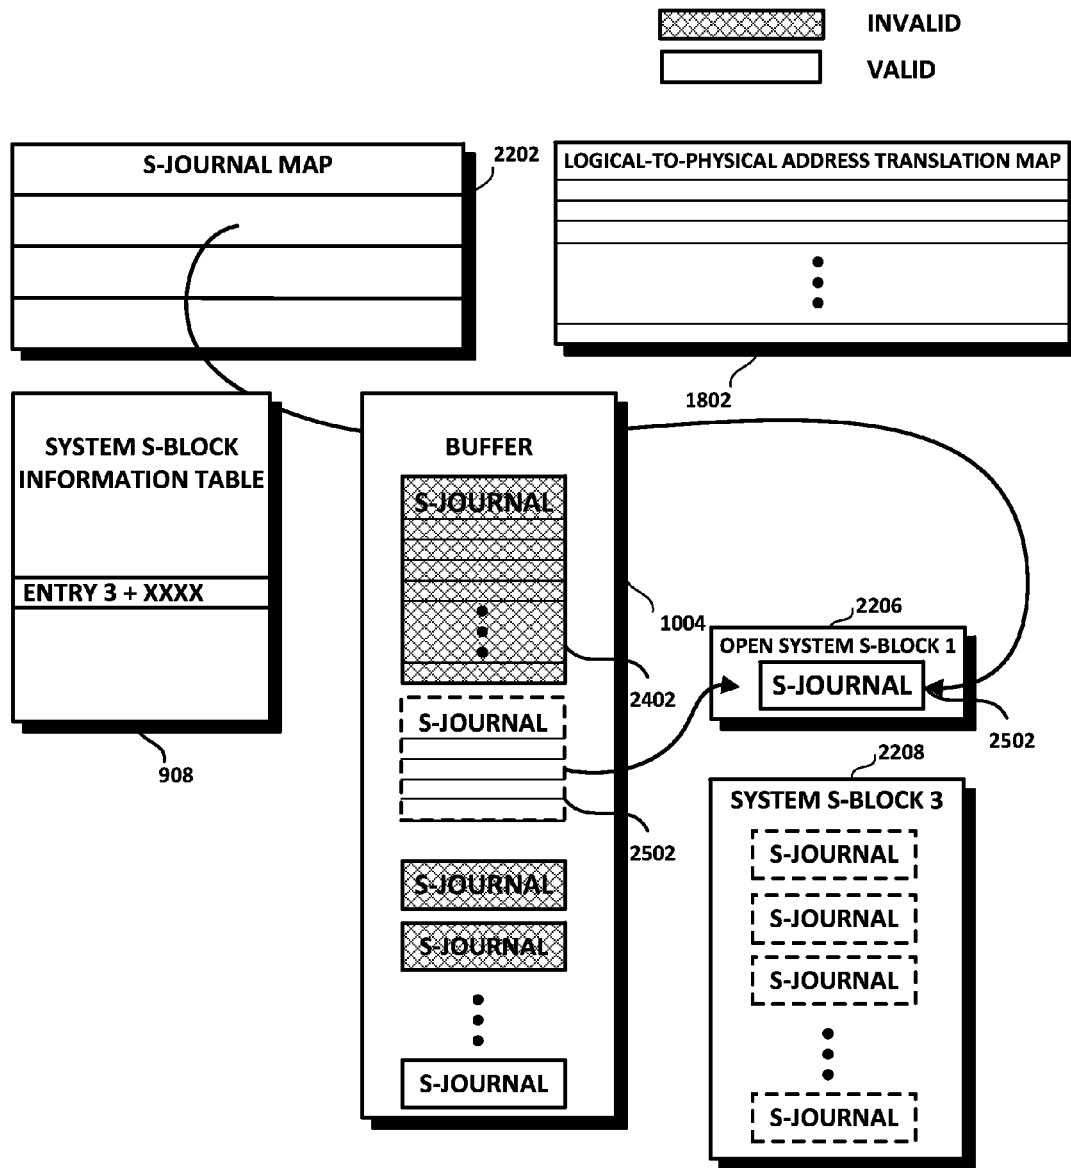


FIG. 26

1

METHODS, DEVICES AND SYSTEMS FOR PHYSICAL-TO-LOGICAL MAPPING IN SOLID STATE DRIVES

BACKGROUND

Due to the nature of flash memory in solid state drives (SSDs), data is typically programmed by pages and erased by blocks. A page in an SSD is typically 8-16 kilobytes (KB) in size and a block consists of a large number of pages (e.g., 256 or 512). Thus, a particular physical location in an SSD (e.g., a page) cannot be directly overwritten without overwriting data in pages within the same block, as is possible in a magnetic hard disk drive. As such, address indirection is needed. Conventional data storage device controllers, which manage the flash memory on data storage devices such as SSDs and interface with the host system, use a Logical to Physical (L2P) mapping system known as Logical Block Addressing (LBA) that is part of the flash translation layer (FTL). When new data comes in replacing older data already written, the data storage device controller causes the new data to be written in a new location and update the logical mapping to point to the new physical location. Since the old physical location no longer holds valid data, it will eventually need to be erased before it can be written again.

Conventionally, a large L2P map table maps logical entries to physical address locations on an SSD. This large L2P map table, which may reside in a volatile memory such as dynamic random access memory (DRAM), is usually updated as writes come in, and saved to non-volatile memory in small sections. For example, if random writing occurs, although the system may have to update only one entry, it may nonetheless have to save to the non-volatile memory the entire table or a portion thereof, including entries that have not been updated, which is inherently inefficient.

FIG. 1 shows aspects of a conventional Logical Block Addressing (LBA) scheme for an SSD. As shown therein, a map table **104** contains one entry for every logical block **102** defined for the data storage device's flash memory **106**. For example, a 64 GB SSD that supports 512 byte logical blocks may present itself to the host as having 125,000,000 logical blocks. One entry in the map table **104** contains the current location of each of the 125,000,000 logical blocks in the flash memory **106**. In a conventional SSD, a flash page holds an integer number of logical blocks (i.e., a logical block does not span across flash pages). In this conventional example, an 8 KB flash page would hold 16 logical blocks (of size 512 bytes). Therefore, each entry in the logical-to-physical map table **104** contains a field **108** identifying the flash die on which the logical block is stored, a field **110** identifying the flash block on which the logical block is stored, another field **112** identifying the flash page within the flash block and a field **114** identifying the offset within the flash page that identifies where the logical block data begins in the identified flash page. The large size of the map table **104** prevents the table from being held inside the SSD controller. Conventionally, the large map table **104** is held in an external DRAM connected to the SSD controller. As the map table **104** is stored in volatile DRAM, it must be restored when the SSD powers up, which can take a long time, due to the large size of the table.

When a logical block is read, the corresponding entry in the map table **104** is read to determine the location in flash memory to be read. A read is then performed to the flash page specified in the corresponding entry in the map table **104**. When the read data is available for the flash page, the data at the offset specified by the map entry is transferred from the

2

SSD to the host. When a logical block is written, the corresponding entry in the map table **104** is updated to reflect the new location of the logical block. It is to be noted that when a logical block is written, the flash memory will initially contain at least two versions of the logical block; namely, the valid, most recently written version (pointed to by the map table **104**) and at least one other, older version thereof that is stale and is no longer pointed to by any entry in the map table **104**. These "stale" data are referred to as garbage, which occupies space that must be accounted for, collected, erased and made available for future use.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows aspects of a conventional Logical Block Addressing (LBA) scheme for SSDs.

FIG. 2 is a diagram showing aspects of the physical and logical data organization of a data storage device according to one embodiment.

FIG. 3 shows a logical-to-physical address translation map and illustrative entries thereof, according to one embodiment.

FIG. 4 shows aspects of a method for updating a logical-to-physical address translation map and for creating an S-Journal entry, according to one embodiment.

FIG. 5 is a block diagram of an S-Journal, according to one embodiment.

FIG. 6 shows an exemplary organization of one entry of an S-Journal, according to one embodiment.

FIG. 7 is a block diagram of a superblock (S-Block), according to one embodiment.

FIG. 8 shows another view of a super page (S-Page), according to one embodiment.

FIG. 9A shows relationships among the logical-to-physical address translation map, S-Journals and S-Blocks, according to one embodiment.

FIG. 9B is a block diagram of an S-Journal Map, according to one embodiment.

FIG. 10 is a block diagram illustrating aspects of a method of updating the logical-to-physical address translation map, according to one embodiment.

FIG. 11 is a block diagram illustrating further aspects of a method of updating the logical-to-physical address translation map, according to one embodiment.

FIG. 12 is a block diagram illustrating still further aspects of a method of updating the logical-to-physical address translation map, according to one embodiment.

FIG. 13 is a block diagram illustrating yet further aspects of a method of updating the logical-to-physical address translation map, according to one embodiment.

FIG. 14 is a block diagram illustrating aspects of a method of updating S-Journals and the S-Journal Map, according to one embodiment.

FIG. 15 is a block diagram illustrating further aspects of a method of updating S-Journals and the S-Journal Map, according to one embodiment.

FIG. 16 is a block diagram illustrating still further aspects of a method of updating S-Journals and the S-Journal Map, according to one embodiment.

FIG. 17 is a block diagram illustrating yet further aspects of a method of updating S-Journals and the S-Journal Map, according to one embodiment.

FIG. 18 is a block diagram illustrating aspects of garbage collection, according to one embodiment.

FIG. 19 is a block diagram illustrating further aspects of garbage collection, according to one embodiment.

FIG. 20 is a block diagram illustrating still further aspects of garbage collection, according to one embodiment.

FIG. 21 is a block diagram illustrating yet further aspects of garbage collection, according to one embodiment.

FIG. 22 is a block diagram illustrating aspects of garbage collecting a system block, according to one embodiment.

FIG. 23 is a block diagram illustrating further aspects of garbage collecting a system block, according to one embodiment.

FIG. 24 is a block diagram illustrating still further aspects of garbage collecting a system block, according to one embodiment.

FIG. 25 is a block diagram illustrating other aspects of garbage collecting a system block, according to one embodiment.

FIG. 26 is a block diagram illustrating still other aspects of garbage collecting a system block, according to one embodiment.

DETAILED DESCRIPTION

System Overview

FIG. 2 is a diagram showing aspects of the physical and logical data organization of a data storage device according to one embodiment. In one embodiment, the data storage device is an SSD. In another embodiment, the data storage device is a hybrid drive including flash memory and rotating magnetic storage media. The disclosure is applicable to both SSD and hybrid implementations, but for the sake of simplicity the various embodiments are described with reference to SSD-based implementations. A data storage device controller 202 according to one embodiment may be configured to be coupled to a host, as shown at reference numeral 218. The host 218 may utilize a logical block addressing (LBA) scheme. While the LBA size is normally fixed, the host can vary the size of the LBA dynamically. For example, the LBA size may vary by interface and interface mode. Indeed, while 512 bytes is most common, 4 KB is also becoming more common, as are 512+ (520, 528 etc.) and 4 KB+ (4 KB+8, 4 KB+16 etc.) formats. As shown therein, the data storage device controller 202 may comprise or be coupled to a page register 204. The page register 204 may be configured to enable the controller 202 to read data from and store data to the data storage device. The controller 202 may be configured to program and read data from an array of flash memory devices responsive to data access commands from the host 218. While the description herein refers to flash memory generally, it is understood that the array of memory devices may comprise one or more of various types of non-volatile memory devices such as flash integrated circuits, Chalco-genide RAM (C-RAM), Phase Change Memory (PC-RAM or PRAM), Programmable Metallization Cell RAM (PMC-RAM or PMCm), Ovonic Unified Memory (OUM), Resistance RAM (RRAM), NAND memory (e.g., single-level cell (SLC) memory, multi-level cell (MLC) memory, or any combination thereof), NOR memory, EEPROM, Ferroelectric Memory (FeRAM), Magnetoresistive RAM (MRAM), other discrete NVM (non-volatile memory) chips, or any combination thereof.

The page register 204 may be configured to enable the controller 202 to read data from and store data to the array. According to one embodiment, the array of flash memory devices may comprise a plurality of non-volatile memory devices in die (e.g., 128 dies), each of which comprises a plurality of blocks, such as shown at 206 in FIG. 2. Other page registers 204 (not shown), may be coupled to blocks on other die. A combination of flash blocks, grouped together, may be called a Superblock or S-Block. In some embodiments, the individual blocks that form an S-Block may be chosen from

one or more dies, planes or other levels of granularity. An S-Block, therefore, may comprise a plurality of flash blocks, spread across one or more die, that are combined together. In this manner, the S-Block may form a unit on which the Flash Management System (FMS) operates. In some embodiments, the individual blocks that form an S-Block may be chosen according to a different granularity than at the die level, such as the case when the memory devices include dies that are sub-divided into structures such as planes (i.e., blocks may be taken from individual planes). According to one embodiment, allocation, erasure and garbage collection may be carried out at the S-Block level. In other embodiments, the FMS may perform data operations according to other logical groupings such as pages, blocks, planes, dies, etc.

In turn, each of the flash blocks 206 comprises a plurality of flash pages (F-Pages) 208. Each F-Page may be of a fixed size such as, for example, 16 KB. The F-Page, according to one embodiment, is the size of the minimum unit of program for a given flash device. As also shown in FIG. 2, each F-Page 208 may be configured to accommodate a plurality of physical pages, hereinafter referred to as E-Pages 210. The term "E-Page" refers to a data structure stored in flash memory on which an error correcting code (ECC) has been applied. According to one embodiment, the E-Page 210 may form the basis for physical addressing within the data storage device and may constitute the minimum unit of flash read data transfer. The E-Page 210, therefore, may be (but need not be) of a predetermined fixed size (such as 2 KB, for example) and determine the size of the payload (e.g., host data) of the ECC system. According to one embodiment, each F-Page 208 may be configured to fit a predetermined plurality of E-Pages 210 within its boundaries. For example, given 16 KB size F-Pages 208 and a fixed size of 2 KB per E-Page 210, eight E-Pages 210 fit within a single F-Page 208, as shown in FIG. 2. In any event, according to one embodiment, a power of 2 multiple of E-Pages 210, including ECC, may be configured to fit into an F-Page 208. Each E-Page 210 may comprise a data portion 214 and, depending on where the E-Page 210 is located, may also comprise an ECC portion 216. Neither the data portion 214 nor the ECC portion 216 need be fixed in size. The address of an E-Page uniquely identifies the location of the E-Page within the flash memory. For example, the E-Page's address may specify the flash channel, a particular die within the identified flash channel, a particular block within the die, a particular F-Page and, finally, the E-Page within the identified F-Page.

To bridge between physical addressing on the data storage device and logical block addressing by the host, a logical page (L-Page) construct is introduced. An L-Page, denoted in FIG. 2 at reference numeral 212 may comprise the minimum unit of address translation used by the FMS. Each L-Page, according to one embodiment, may be associated with an L-Page number. The L-Page numbers of L-Pages 212, therefore, may be configured to enable the controller 202 to logically reference host data stored in one or more of the physical pages, such as the E-Pages 210. The L-Page 212 may also be utilized as the basic unit of compression. According to one embodiment, unlike F-Pages 208 and E-Pages 210, L-Pages 212 are not fixed in size and may vary in size, due to variability in the compression of data to be stored. Since the compressibility of data varies, a 4 KB amount of data of one type may be compressed into a 2 KB L-Page while a 4 KB amount of data of a different type may be compressed into a 1 KB L-Page, for example. Due to such compression, therefore, the size of L-Pages may vary within a range defined by a minimum compressed size of, for example, 24 bytes to a maximum uncompressed size of, for example, 4 KB or 4 KB+. Other

5

sizes and ranges may be implemented. As shown in FIG. 2, L-Pages 212 need not be aligned with the boundaries of E-Page 210. Indeed, L-Pages 212 may be configured to have a starting address that is aligned with an F-Page 208 and/or E-Page 210 boundary, but also may be configured to be unaligned with either of the boundaries of an F-Page 208 or E-Page 210. That is, an L-Page starting address may be located at a non-zero offset from either the start or ending addresses of the F-Pages 208 or the start or ending addresses of the E-Pages 210, as shown in FIG. 2. As the L-Pages 212 are not fixed in size and may be smaller than the fixed-size E-Pages 210, more than one L-Page 212 may fit within a single E-Page 210. Similarly, as the L-Pages 212 may be larger in size than the E-Pages 210, L-Pages 212 may span more than one E-Page, and may even cross the boundaries of F-Pages 210, shown in FIG. 2 at numeral 217.

For example, where the LBA size is 512 or 512+ bytes, a maximum of, for example, eight sequential LBAs may be packed into a 4 KB L-Page 212, given that an uncompressed L-Page 212 may be 4 KB to 4 KB+. It is to be noted that, according to one embodiment, the exact logical size of an L-Page 212 is unimportant as, after compression, the physical size may span from few bytes at minimum size to thousands of bytes at full size. For example, for 4 TB SSD device, 30 bits of addressing may be used to address each L-Page 212 to cover for an amount of L-Pages that could potentially be present in such a SSD.

FIG. 3 shows a logical-to-physical address translation map and illustrative entries thereof, according to one embodiment. As the host data is referenced by the host in L-Pages 212 and as the data storage device stores the L-Pages 212 in one or more contiguous E-Pages 210, a logical-to-physical address translation map is required to enable the controller 202 to associate an L-Page number of an L-Page 212 to one or more E-Pages 210. Such a logical-to-physical address translation map is shown in FIG. 3 at 302 and, in one embodiment, is a linear array having one entry per L-Page 212. Such a logical-to-physical address translation map 302 may be stored in a volatile memory, such as a DRAM or SRAM. FIG. 3 also shows the entries in the logical-to-physical address translation map for four different L-Pages 212, which L-Pages 212 in FIG. 3 are associated with L-Page numbers denoted as L-Page 1, L-Page 2, L-Page 3 and L-Page 4. According to one embodiment, each L-Page stored in the data storage device may be pointed to by a single and unique entry in the logical-to-physical address translation map 302. Accordingly, in the example being developed herewith, four entries are shown. As shown at 302, each entry in the map 302 may comprise an L-Page number, which may comprise an identification of the physical page (e.g., E-Page) containing the start address of the L-Page being referenced, the offset of the start address within the physical page (e.g., E-Page) and the length of the L-Page. In addition, a plurality of ECC bits may provide error correction functionality for the map entry. For example, and as shown in FIG. 3, and assuming an E-Page size of 2 KB, L-Page 1 may be referenced in the logical-to-physical address translation map 302 as follows: E-Page 1003, offset 800, length 1624, followed by a predetermined number of ECC bits (not shown). That is, in physical address terms, the start of L-Page 1 is within (not aligned with) E-Page 1003, and is located at an offset from the starting physical location of the E-Page 1003 that is equal to 800 bytes. Compressed L-Page 1, furthermore, extends 1,624 bytes, thereby crossing an E-Page boundary to E-Page 1004. Therefore, E-Pages 1003 and 1004 each store a portion of the L-Page 212 denoted by L-Page number L-Page 1. Similarly, the compressed L-Page referenced by L-Page number L-Page 2 is stored entirely within

6

E-Page 1004, and begins at an offset therein of 400 bytes and extends only 696 bytes within E-Page 1004. The compressed L-Page associated with L-Page number L-Page 3 starts within E-Page 1004 at an offset of 1,120 bytes (just 24 bytes away from the boundary of L-Page 2) and extends 4,096 bytes past E-Page 1005 and into E-Page 1006. Therefore, the L-Page associated with L-Page number L-Page 3 spans a portion of E-Page 1004, all of E-Page 1005 and a portion of E-Page 1006. Finally, the L-Page associated with L-Page number L-Page 4 begins within E-Page 1006 at an offset of 1,144 bytes, and extends 3,128 bytes to fully span E-Page 1007, crossing an F-Page boundary into E-Page 1008 of the next F-Page.

Collectively, each of these constituent identifier fields (E-Page, offset, length and ECC) making up each entry of the logical-to-physical address translation map 302 may be, for example, 8 bytes in size. That is, for an exemplary 4 TB drive, the address of the E-Page may be 32 bits in size, the offset may be 12 bits (for E-Page data portions up to 4 KB) in size, the length may be 10 bits in size and the ECC field may be provided. Other organizations and bit-widths are possible. Such an 8 byte entry may be created each time an L-Page is written or modified, to enable the controller 202 to keep track of the host data, written in L-Pages, within the flash storage. This 8-byte entry in the logical-to-physical address translation map may be indexed by an L-Page number of LPN. In other words, according to one embodiment, the L-Page number functions as an index into the logical-to-physical address translation map 302. It is to be noted that, in the case of a 4 KB sector size, the LBA is the same as the LPN. The LPN, therefore, may constitute the address of the entry within the volatile memory. When the controller 202 receives a read command from the host 218, the LPN may be derived from the supplied LBA and used to index into the logical-to-physical address translation map 302 to extract the location of the data to be read in the flash memory. When the controller 202 receives a write command from the host, the LPN may be constructed from the LBA and the logical-to-physical address translation map 302 may be modified. For example, a new entry therein may be created. Depending upon the size of the volatile memory storing the logical-to-physical address translation map 302, the LPN may be stored in a single entry or broken into, for example, a first entry identifying the E-Page containing the starting address of the L-Page in question (plus ECC bits) and a second entry identifying the offset and length (plus ECC bits). According to one embodiment, therefore, these two entries may together correspond and point to a single L-Page within the flash memory. In other embodiments, the specific format of the logical-to-physical address translation map entries may be different from the examples shown above.

S-Journals and S-Journal Map

As the logical-to-physical address translation map 302 may be stored in a volatile memory, it may need to be rebuilt upon startup or any other loss of power to the volatile memory. This, therefore, requires some mechanism and information to be stored in a non-volatile memory that will enable the controller 202 to reconstruct the logical-to-physical address translation map 302 before the controller can “know” where the L-Pages are stored in the non-volatile memory after startup or after a power-fail event. According to one embodiment, such mechanism and information are embodied in a construct that may be called a System Journal, or S-Journal. According to one embodiment, the controller 202 may be configured to maintain, in the plurality of non-volatile memory devices (e.g., in one or more of the blocks 206 in one or more die, channel or plane), a plurality of

7

S-Journals defining physical-to-logical address correspondences. According to one embodiment, each S-Journal covers a pre-determined range of physical pages (e.g., E-Pages). According to one embodiment, each S-Journal may comprise a plurality of journal entries, with each entry being configured to associate one or more physical pages, such as E-Pages, to the L-Page number of each L-Page. According to one embodiment, each time the controller 202 restarts or whenever the logical-to-physical address translation map 302 is to be rebuilt either partially or entirely, the controller 202 reads the S-Journals and, from the information read from the S-Journal entries, rebuilds the logical-to-physical address translation map 302.

FIG. 4 shows aspects of a method for updating a logical-to-physical address translation map and for creating an S-Journal entry, according to one embodiment. As shown therein, to ensure that the logical-to-physical address translation map 302 is kept up-to-date, whenever an L-Page is written or otherwise updated as shown at block B41, the logical-to-physical address translation map 302 may be updated as shown at B42. As shown at B43, an S-Journal entry may also be created, storing therein information pointing to the location of the updated L-Page. In this manner, both the logical-to-physical address translation map 302 and the S-Journals are updated when new writes occur (e.g., as the host issues writes to non-volatile memory, as garbage collection/wear leveling occurs, etc.). Write operations to the non-volatile memory devices to maintain a power-safe copy of address translation data may be configured, therefore, to be triggered by newly created journal entries (which may be just a few bytes in size) instead of re-saving all or a portion of the logical-to-physical address translation map, such that Write Amplification (WA) is reduced. The updating of the S-Journals ensures that the controller 202 can access a newly updated L-Page and that the logical-to-physical address translation map 302 may be reconstructed upon restart or other information-erasing power event affecting the volatile memory in which the logical-to-physical address translation map is stored. Moreover, in addition to their utility in rebuilding the logical-to-physical address translation map 302, the S-Journals are useful in enabling effective Garbage Collection (GC). Indeed, the S-Journals may contain the last-in-time update to all L-Page numbers, and also may contain stale entries, entries that do not point to a valid L-Page.

According to one embodiment, the S-Journals may be the main flash management data written to the non-volatile memory. S-Journals may contain mapping information for a given S-Block and may contain the Physical-to-Logical (P2L) information for a given S-Block. FIG. 5 is a block diagram showing aspects of an S-Journal, according to one embodiment. As shown therein, each S-Journal 502 covers a predetermined physical region of the non-volatile memory such as, for example, 32 E-Pages as shown at 506, which are addressable using 5 bits. Each S-Journal 502 may be identified by an S-Journal Number, which may be part of a header 504 that could include other information about the S-Journal. The S-Journal Number may comprise a portion of the address of the first physical page covered by the S-Journal. For example, the S-Journal Number of S-Journal 502 may comprise, for example, the 27 Most Significant Bits (MSB) of the first E-Page address covered by this S-Journal 502.

FIG. 6 shows an exemplary organization of one entry 602 of an S-Journal 502, according to one embodiment. Each entry 602 of the S-Journal 502 may point to the starting address of one L-Page, which is physically addressed in E-Pages. Each entry 602 may comprise, for example, a number (5, for example) of Least Significant Bits (LSBs) of the

8

address of the E-Page containing the start L-Page. The full E-Page address is obtained by concatenating these 5 LSBs with the 27 MSBs of the S-Journal Number in the header 504. In addition, the entry 602 may comprise the L-Page number, its offset within the identified E-Page and its size. For example, each entry 602 of an S-Journal may comprise the 5 LSBs of the address of first E-Page covered by this S-Journal entry, 30 bits of L-Page number, 9 bits of E-Page offset and 10 bits of L-Page size, adding up to an overall size of about 7 bytes. Various other internal journal entry formats may be used in other embodiments.

According to one embodiment, due to the variability in the compression or the host configuration of the data stored in L-Pages, a variable number of L-Pages may be stored in a physical area, such as a physical area equal to 32 E-Pages, as shown at 506. As a result of the use of compression and the consequent variability in the sizes of L-Pages, S-Journals may comprise a variable number of entries. For example, according to one embodiment, at maximum compression, an L-Page may be 24 bytes in size and an S-Journal may comprise over 2,500 entries, referencing an equal number of L-Pages, one L-Page per S-Journal entry 602.

As noted above, an S-Journal may be configured to contain mapping information for a given S-Block. More precisely, according to one embodiment, S-Journals contain the mapping information for a predetermined range of E-Pages within a given S-Block. FIG. 7 is a block diagram of a S-Block, according to one embodiment. As shown therein, an S-Block 702 may comprise one flash block (F-Block) 704 (as also shown at 206 in FIG. 2) per die. An S-Block, therefore, may be thought of as a collection of F-Blocks, one F-Block per die, that are combined together to form a unit of the Flash Management System. According to one embodiment, allocation, erasure and GC may be managed at the S-Block level. Each F-Block 704, as shown in FIG. 7, may comprise a plurality of flash pages (F-Page) such as, for example, 256 or 512 F-Pages. An F-Page, according to one embodiment, may be the size of the minimum unit of program for a given non-volatile memory device. FIG. 8 shows a super page (S-Page), according to one embodiment. As shown therein, an S-Page 802 may comprise one F-Page per F-Block of an S-Block, meaning that an S-Page spans across an entire S-Block.

Relationships Among Various Data Structures

FIG. 9A shows relationships among the logical-to-physical address translation map, the S-Journal map and S-Blocks, according to one embodiment. Reference 902 denotes an entry in the logical-to-physical address translation map (stored in DRAM in one embodiment). According to one embodiment, the logical-to-physical address translation map may be indexed by L-Page number, in that there may be one entry 902 per L-Page in the logical-to-physical address translation map. The physical address of the start of the L-Page in the flash memory and the size thereof may be given in the map entry 902; namely by E-Page address, offset within the E-Page and the size of the L-Page. As noted earlier, the L-Page, depending upon its size, may span one or more E-Pages and may span F-Pages and F-Blocks as well.

As shown at 904, the volatile memory (e.g., DRAM) may also store a System Journal (S-Journal) map. An entry 904 in the S-Journal map stores information related to where an S-Journal is physically located in the non-volatile memory. For example, the 27 MSBs of the E-Page physical address where the start of the L-Page is stored may constitute the S-Journal Number (as previously shown in FIG. 5). The S-Journal map entry 904 in the volatile memory may also include the address of the S-Journal in non-volatile memory, referenced in system E-Pages. From the S-Journal map entry

904 in volatile memory, System S-Block Information 908 may be extracted. The System S-Block Information 908 may be indexed by System S-Block (S-Block in the System Band) and may comprise, among other information regarding the S-Block, the size of any free or used space in the System S-Block. Also from the S-Journal map entry 904, the physical location (expressed in terms of E-Pages in the System Band) of the referenced S-Journal in non-volatile memory 910 may be extracted.

The System Band, according to one embodiment, does not contain L-Page data and may contain File Management System (FMS) meta-data and information. The System Band may be configured as lower page only for reliability and power fail simplification. During normal operation, the System Band need not be read except during Garbage Collection. The System Band may be provided with significantly higher overprovisioning than the data band for overall WA optimization. Other bands include the Hot Band, which may contain L-Page data and is frequently updated, and the Cold Band, which may be less frequently updated and may comprise more static data, such as data that may have been collected as a result of GC. According to one embodiment, the System, Hot and Cold Bands may be allocated by an S-Block basis.

As noted above, each of these S-Journals in non-volatile memory may comprise a collection of S-Journal entries and cover, for example, 32 E-Pages worth of data. These S-Journals in non-volatile memory 910 enable the controller 202 to rebuild not only the logical-to-physical address translation map in volatile memory, but also the S-Journal map, the User S-Block Information 906, and the System S-Block Information 908, in volatile memory.

FIG. 9B is a block diagram of an S-Journal Map 912, according to one embodiment. The S-Journal Map 912 may be indexed by S-Block number and each entry thereof may point to the start of the first S-Journal for that S-Block which, in turn, may cover a predetermined number of E-Pages (e.g., 32) of that S-Block. The controller 202 may be further configured to build or rebuild a map of the S-Journals and store the resulting S-Journal Map in volatile memory. That is, upon restart or upon the occurrence of another event in which power fails or after a restart subsequent to error recovery, the controller 202 may read the plurality of S-Journals in a predetermined sequential order, build a map of the S-Journals stored in the non-volatile memory devices based upon the sequentially read plurality of S-Journals, and store the built S-Journal Map 912 in the volatile memory.

Updating the Logical-to-Physical Address Translation Map

FIGS. 10-13 are block diagrams illustrating aspects of a method of updating the logical-to-physical address translation map, according to one embodiment. As shown in FIG. 10, a logical-to-physical address translation map 1002 contains an entry (e.g., a location of an L-Page) for L-Page 100, which has a length of 3,012 bytes. In this example, L-Page 100 is stored in S-Block 15, as shown at 1006. A buffer 1004 (such as a static random access memory (SRAM)) in or coupled to the controller 202 may store the S-Journal that contains the P2L information for S-Block 15 in which the L-Page 100 resides. What is shown in the buffer 1004 may actually reside in DRAM in some embodiments. The User S-Block Info 906, whose entries are indexed by S-Block, may comprise for each S-Block, among other information regarding the S-Block, the (e.g., exact or approximate) size of the free or used space in the S-Block. As shown in FIG. 10, the entry in the User S-Block Info 906 for S-Block 15 is shown. FIG. 10 shows an illustrative state of these constituent functional blocks before an update to L-Page 100 is processed by the controller 202.

As shown in FIG. 11 at 1102, an updated L-Page 100 is received, with a new length of 1,534 bytes. Responsive to the receipt of the updated L-Page 100, the L-Page 100 information may be fetched from the logical-to-physical address translation map 1002 and the length (3,012 bytes) of the (now obsolete) L-Page 100 may be extracted therefrom and used to correspondingly and precisely increase the tracked free space value of S-Block 15. In particular, the User S-Block Information 906 may be updated with data indicating that S-Block 15 now has 3,012 additional bytes of free space, now that the original data for L-Page 100 is now stale (e.g., obsolete).

As shown in FIG. 12, the logical-to-physical address translation map 1002 may be updated to accommodate the updated L-Page. For example, the length information is now 1,534 bytes. Thereafter, the S-Journal in the buffer 1004 for the particular portion of S-Block 15 where the update occurs may now be updated with the new L-Page information, including length information and the L-Page newly received from, for example, a compressor may be read into the buffer 1004. While the new L-Page is in the buffer, the entry for L-Page 100 in the logical-to-physical address translation map 1002 may temporarily reflect its location in the buffer, as shown by the arrow labeled "1". Later, the updated L-Page 100 may be flushed to the Hot S-Block 1008, at the E-Page address specified by the newly-created entry in the now updated S-Journal still in the buffer 1004. The mapping table entry in the logical-to-physical address translation map 1002 is updated to reflect the physical E-Page address and offset of the L-Page's final destination, as suggested by arrow "2".

As shown in FIG. 13, at some later point in time such as, for example, after accumulating a sufficient number of new entries, the S-Journal in the volatile memory buffer 1004 may be written out to non-volatile memory, such as to the System S-Block 1010. The System S-Block 1010 may be a portion of the flash memory allocated by the controller firmware to store S-Journals. The saving of the S-Journal in non-volatile memory enables the later reconstruction of the logical-to-physical address translation map, as needed.

Updating the S-Journals and S-Journal Map

FIGS. 14-17 are block diagrams illustrating aspects of a method of updating S-Journals and the S-Journal Map in the System Band, according to one embodiment. The figures illustrate additional details relating to the update mechanisms of S-Journal triggered by the same example update to L-Page 100. As shown therein, an S-Journal Map 1402 in volatile memory (e.g., synchronous dynamic random-access memory (SDRAM)), may contain an entry that points to the physical location, in non-volatile memory, of the S-Journal that contains the P2L mapping data associated with L-Page 100. In the illustrative example, the relevant S-Journal (i.e., S-Journal for a particular portion of S-Block 15 storing L-Page 100) is located in System S-Block 3 (at 1408), starting at E-Page 42.

Thereafter, as shown FIG. 15, L-Page 100 is updated, with a length of 1,534 bytes. The information (e.g., E-Page address, offset and length) regarding current L-Page 100 may then be fetched from the logical-to-physical address translation map 1002. Thereafter, the address of the E-Page storing L-Page 100 may be used to extract the S-Journal Number used to locate the S-Journal entry within the S-Journal Map 1402. For example, in the case wherein, as shown at FIGS. 5 and 6, the S-Journal Number may comprise the 27 MSBs of the first 32-bit E-Page address. This S-Journal number may then be used to index into the S-Journal Map 1402 to obtain the location, which may then be used to identify the System S-Block that contains the S-Journal of interest. Thereafter, as shown at FIG. 15, the System S-Block Information 908 may be updated to reflect the fact that the journal entry for L-Page

11

100 in the S-Journal at S-Block 3 is now invalid, as the space previously occupied by that entry has now been de-allocated, in view of the recent update to L-Page 100. That de-allocated space is now free space, and must be accounted for. Alternatively, used space may be accounted for and the free space derived therefrom. The entry in the System S-Block information 908 for S-Block 3 may, therefore, be incremented by a space taken by one S-Journal entry; namely, 7 bytes in the exemplary implementation being developed herein. This empty space may thereafter be taken into account during GC of the System Band.

As shown in FIG. 16, L-Page 100 may then be written to the buffer 1004, as shown at 1005, whereupon the logical-to-physical address translation map 1002 may be updated to point to the buffer 1004 (indicating the physical location where L-Page 100 is stored) and to store the length of the updated L-Page 100 (1,534 in the example being developed herein). It is to be noted that, at this point, System S-Block 3 still includes an S-Journal entry that points to S-Block 15 at reference 1006 as containing L-Page 100. However, the System S-Block Info for S-Block 3 has been updated so as to indicate that the space presently occupied by the now outdated S-Journal entry within the System S-Block 3 at reference 1408 is, in fact, free space.

FIG. 16 also shows a different S-Journal in the buffer. In one embodiment, the S-Journal in the buffer 1004 is for recording new L-Pages written to the Open Hot S-Block 7. In this example, the S-Journal in the buffer 1004 (which has not yet been flushed to the non-volatile memory) is updated with the L-Page number of updated L-Page 100 (now in the buffer also) and its length information, as indicated by the arrow at 1007 in FIG. 16. L-Page 100 may now be flushed from the buffer 1004 to the currently open Hot S-Block 7, as referenced at 1404 in FIG. 16. At that point, the S-Journal in the buffer is updated to reflect that updated L-Page 100 has been written to the non-volatile memory. Alternatively, the logical-to-physical address translation map 1002 and the S-Journal may be updated with the address of the updated L-Page 100 prior to the updated L-Page 100 being written to the non-volatile memory. In one embodiment, an updated entry in that S-Journal will contain the E-Page pointed to by the new L-Page 100's entry in the logical-to-physical address translation map 1002 and additional information about L-Page 100, as previously shown in FIG. 6.

As shown in FIG. 17, after accumulating sufficient data, the S-Journal in the buffer 1004 may be written out to an open System S-Block 1 (referenced at 1406). In this example, this S-Journal is written out to E-Page 19 of open System S-Block 1. The S-Journal Map 1402 now comprises one entry indicating that a portion of S-Block 15 is stored in System S-Block 3, E-Page 42. However, the entry within that S-Journal which pointed to the old location of L-Page 100 within S-Block 15 is no longer valid. The S-Journal Map 1402 also comprises one entry indicating that S-Journal for S-Block 7 is stored in System S-Block 1, E-Page 19. That S-Journal includes a valid entry that points to open Hot S-Block 7 at 1404 as the location in non-volatile memory of the updated L-Page 100. In this manner, the S-Journals containing the P2L information are updated in the buffer 1004, and are eventually written out to non-volatile memory. In addition, the S-Journal Map 1402 may be suitably updated to point to such newly updated S-Journals. Moreover, the logical-to-physical address translation map 1002 may be suitably updated to point to the correct open Hot S-Block.

Advantageously, on a new write, the S-Journal construct as shown and described herein minimizes write overhead. According to one embodiment, a new write operation neces-

12

sitates writing a new S-Journal entry in the non-volatile memory. Since there is only a small amount of S-Journal entry data per L-Page generated (e.g., 7 bytes as described herein), the WA due to system data writes is reduced as compared to conventional systems. S-Journal data is effectively a massive command history. Updated S-Journal data is pooled up and written out sequentially to the System Band. The S-Journal system shown and described herein is efficient, as the mapping information that is generated is the same as the data to be written, with no additional unchanged entries, as only that which is changed is written to non-volatile (e.g., flash) memory, as opposed to all or a portion of a logical-to-physical map.

On power up, all the S-Journal data in the System Band may be read and processed into DRAM to rebuild the logical-to-physical address translation map in volatile memory. According to one embodiment, this is done in the order in which the S-Blocks are allocated to the System Band. In one embodiment, hardware support is used to enable this large data load to occur quickly to meet power-up timing requirement (this may not be practical without hardware, since this is essentially generating the logical-to-physical address translation map from the command history). According to one embodiment, an S-Journal Map is also constructed at power-up to point to valid S-Journals stored in the System Band.

Garbage Collection of User Data S-Blocks

FIGS. 18-21 are block diagrams illustrating aspects of garbage collection, according to one embodiment. As shown therein, the data in the User S-Block Information 906 may be scanned to select the "best" S-Block to garbage collect. There are a number of criteria that may be evaluated to select which S-Block to garbage collect. For example, the best S-Block to garbage collect may be that S-Block having the largest amount of free space and the lowest Program Erase (PE) count. Alternatively, these and/or other criteria may be weighted to select the S-Block to be garbage collected. For purposes of example, the S-Block selected to be garbage collected in FIGS. 18-21 is S-Block 15, which is referenced by information entry 15 within the User S-Block Information 906, showing some amount+3,012 bytes of free space (the additional 3,012 bytes to account for the recently obsoleted L-Page 100). It is to be noted that the User S-Block Information 906 may comprise, among other items of information, a running count of the number of PE cycles undergone by each tracked S-Block, which may be evaluated in deciding which S-Block to garbage collect. As shown at 1006, S-Block 15 has a mix of valid data (hashed blocks) and invalid data (non-hashed blocks).

Now that S-Block 15 has been selected for GC, the S-Journal Map (see 912 in FIG. 9B) may be consulted (e.g., indexed into by the S-Block number) to find the location in non-volatile memory (e.g., E-Page address) of the corresponding S-Journals for that S-Block. To illustrate an example, one S-Journal pointed to by the S-Journal Map 912 is located using the S-Journal Number (27 MSb of E-Page Address), and read into the buffer 1004, as shown in FIG. 18. That is, the 1 or more E-Pages in the System S-Block 1804 pointed to by S-Journal Map 912 are accessed and the S-Journal stored beginning at that location may be read into the buffer 1004. In one embodiment, the S-Journal map also contains the length of the S-Journal since an S-Journal may span one or more E-Pages. An S-Journal may be quite large and, therefore, may be read in pieces and processed as available.

Thereafter, each P2L entry in the S-Journal in the buffer 1004 may then be compared to the corresponding entry in the logical-to-physical address translation map 1802. For each entry in the S-Journal in the buffer 1004, it may be determined

whether the physical address for the L-Page of that entry matches the physical address of the same L-Page in the corresponding entry in the logical-to-physical address translation map **1802**. If the two match, that entry in the S-Journal is valid. Conversely, if the address for the L-Page in the S-Journal does not match the entry for that L-Page in the logical-to-physical address translation map, that entry in the S-Journal is not valid.

According to one embodiment, as valid entries are found in the S-Journal whose entries are being parsed and compared, the referenced L-Pages may be read out of S-Block **15** and written to the buffer **1004**, as shown in FIG. **19**. The same process may be used for other S-Journals covering S-Block **15** until the entire S-Block is processed. As also shown in FIG. **19** at reference **1006**, S-Block **15** now contains only invalid data. This is because the data indicated as valid by entries in S-Journals for S-Block **15** have been preserved and will soon be moved to a new S-Block. In the illustrated example, as the entries in the example S-Journal of S-Block **15** in System S-Block **1804** point to such invalid data, that S-Journal is shown as being hashed, indicating that it is now stale. The logical-to-physical address translation map **1802** may then be updated, generating a new E-Page starting address for the valid data read into the buffer **1004**. It is to be noted that during the update of the logical-to-physical address translation map, the map may be rechecked for valid entries and may be locked during the map update process to guarantee atomicity. The valid data will also necessitate new S-Journal entries be generated in S-Journal **1005**, which in one embodiment is for the Cold S-Block **1801**.

In one embodiment, the valid data may then be written out to the Cold S-Block **1801** (the Hot S-Block being used for recently written host data, not garbage collected data), as shown at FIG. **20**. At some later time (e.g., after a sufficient number of entries have been populated), S-Journal **1005** may be written out to the System S-Block **1804** in the System Band. S-Block **15** has now been garbage collected and User S-Block Info **906** now indicates that the entire S-Block **15** is free space. S-Block **15** may thereafter be erased, its PE count updated and made available for new writes. It is to be noted that an invalid S-Journal is still present in System S-Block **1804**. The space in flash memory in the System Band occupied by this invalid S-Journal may be garbage collected at some later time.

Garbage Collection on System S-Blocks

FIGS. **22-26** are block diagrams illustrating aspects of garbage collecting a system block, according to one embodiment. In the example being developed in FIGS. **22-26**, it is assumed that System S-Block **3**, shown at referenced **2208** in FIG. **22**, has been picked for garbage collection. Once a System S-Block has been picked, all of the E-pages (and by extension, all S-Journals contained therein) within the picked System S-Block may be read (sequentially or non-sequentially) into the buffer **1004**.

As suggested in FIG. **23**, the S-Journal numbers for one or more of the S-Journals read into the buffer **1004** may then be extracted from the headers of the S-Journals. Each such System S-Journal number may then be used to look up in the S-Journal Map **2202** to determine whether the corresponding S-Journal is still valid. According to one embodiment, invalid S-Journals are those S-Journals whose S-Journal number is not matched by a corresponding entry in the S-Journal Map **2202**, which has the most updated information on where S-Journals are physically stored. For example, if the entry in the S-Journal Map **2202** for S-Journal Number "12345" points to an E-Page within System S-Block **120**, the copy of S-Journal "12345" in S-Block **3** (the S-Block being garbage

collected) is obsolete. Likewise, if the S-Journal Map entry instead points to the E-Page from S-Block **3** where S-Journal "12345" currently resides, S-Journal "12345" is still valid.

In one embodiment, a valid S-Journal being garbage collected may include a mix of valid and invalid entries, thus individual checking of the entries is needed. As shown in FIG. **24**, each entry in each valid S-Journal **2402** may then be matched with a corresponding entry in the logical-to-physical address translation map **1802** in memory. That is, the E-Page address of the L-Page referenced by each S-Journal entry may be compared with the E-Page address of the L-Page specified in the logical-to-physical address translation map **1802**. If the two match, that S-Journal entry is valid. According to one embodiment, it may be necessary to compare the E-Page address and the offset within the E-Page in the S-Journal with the E-Page address and the offset within the E-Page of the L-Page specified in the logical-to-physical address translation map **1802**. Conversely, if the E-page address (or E-page address and offset) for the L-Page in the S-Journal does not match the E-Page address (or E-page address and offset) in the entry for that L-Page in the address translation map **1802**, that entry in the S-Journal is not valid.

According to one embodiment, as valid entries are identified in the S-Journal **2402** (whose entries are being parsed and compared), they may be copied to a new version **2502** of the S-Journal **2402**, as shown in FIG. **25**. The S-Journal **2502**, according to one embodiment, may have the same S-Journal number as that of the S-Journal **2402**. In this manner, each S-Journal loaded into the buffer **1004** from the S-Block picked for GC may be first determined to be valid or invalid at the journal level, and then, if determined valid, compacted to contain only valid entries. According to one embodiment, invalid S-Journal entries are simply not copied to the new version of the S-Journal **2502**. Accordingly, it is expected that the new version **2502** of the S-Journal will be smaller (i.e., comprise fewer entries) than the older version **2402** thereof, presuming that the S-Journal **2402** had one or more obsolete entries. It is to be noted that if an S-Journal has all valid entries, the size of the new version thereof will remain the same.

As shown in FIG. **26**, the S-Journal **2502** (the new version of the S-Journal **2402**, which is now invalid or obsolete) may then be written to the current open System S-Block which, in FIG. **26**, is Open System S-Block **1**, reference numeral **2206**. Thereafter, the S-Journal Map **2202** may be updated with the new location of the S-Journal **2502** in Open System S-Block **1**. At the end of this process, the S-Journals in System S-Block **3** (**2208**) have been garbage collected and System S-Block **3** may then be erased and made available for future programming.

While certain embodiments of the disclosure have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the disclosure. Indeed, the novel methods, devices and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the disclosure. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the disclosure. For example, those skilled in the art will appreciate that in various embodiments, the actual physical and logical structures may differ from those shown in the figures. Depending on the embodiment, certain steps described in the example above may be removed, others may be added. Also, the features and attributes of the specific embodiments disclosed above may be combined in different

15

ways to form additional embodiments, all of which fall within the scope of the present disclosure. Although the present disclosure provides certain preferred embodiments and applications, other embodiments that are apparent to those of ordinary skill in the art, including embodiments which do not provide all of the features and advantages set forth herein, are also within the scope of this disclosure. Accordingly, the scope of the present disclosure is intended to be defined only by reference to the appended claims.

The invention claimed is:

1. A data storage device, comprising:

a plurality of non-volatile memory devices, each configured to store a plurality of physical pages, each of the plurality of physical pages being stored at a predetermined physical location within the plurality of non-volatile devices;

a volatile memory comprising a logical-to-physical address translation map configured to enable a controller to determine a physical location, within one or more physical pages, of the data stored in a plurality of logical pages (L-Pages); and

the controller coupled to the plurality of non-volatile memory devices and configured to program data to and read data from the plurality of non-volatile memory devices, the data being stored in the plurality of logical pages (L-Pages), each of the plurality of L-Pages being associated with an L-Page number that is configured to enable the controller to logically reference data stored in the one or more of the physical pages using the logical-to-physical address translation map, wherein an entry in the logical-to-physical address translation map includes an association between an L-page number and at least one of a length of the L-page or a size of error correcting code bits to be applied to the entry in the logical-to-physical address translation map,

wherein the controller is configured to maintain, in the plurality of non-volatile memory devices, a plurality of journals defining physical-to-logical correspondences, each of the plurality of journals being associated with a journal number, each journal covering a pre-determined range of physical pages and comprising a plurality of journal entries, each entry being configured to associate one or more physical pages to each L-Page, wherein the controller is configured to read the plurality of journals upon startup and rebuild the logical-to-physical address translation map stored in the volatile memory from the read plurality of journals.

2. The data storage device of claim 1, wherein the controller is further configured to, upon an update to one of the plurality of L-Pages, create a new entry in one of the plurality of journals.

3. The data storage device of claim 2, wherein write operations to the non-volatile memory devices to maintain a power-safe copy of translation data are configured to be triggered by newly created journal entries instead of saving at least portions of the logical-to-physical address translation map, such that write amplification is reduced.

4. The data storage device of claim 2, wherein the new entry indicates a physical location, within a physical page, of a start of the updated L-Page.

5. The data storage device of claim 2, wherein the controller is further configured to update a free space accounting by an amount corresponding to a length of the L-Page prior to being updated.

6. The data storage device of claim 2, wherein the controller is further configured to update the logical-to-physical address translation map with the physical location, within one

16

or more physical pages, of the data referenced by the L-Page number of the updated L-Page.

7. The data storage device of claim 1, wherein at least one of the plurality of L-Pages is unaligned with physical page boundaries.

8. The data storage device of claim 1, wherein the physical pages are implemented as Error Correcting Code (ECC) pages (E-Pages) and wherein the plurality of devices comprises a plurality of flash memory blocks, each flash memory block comprising a plurality of flash memory pages (F-Pages), each of the F-Pages comprising a plurality of the E-Pages, each of the plurality of E-Pages being stored at a predetermined physical location within the plurality of devices.

9. The data storage device of claim 1, further comprising a plurality of super blocks (S-Blocks), each comprising one or more flash memory blocks per device and wherein each of the plurality of journal entries is configured to associate one or more of the physical pages of the S-Block to each L-Page.

10. The data storage device of claim 9, wherein the controller is further configured to garbage collect by at least:

selecting one of the plurality of S-Blocks to garbage collect;

comparing each entry in a journal for the selected S-Block to entries in the logical-to-physical address translation map and designating entries that match as valid and entries that do not match as invalid;

reading the L-Pages corresponding to the valid entries;

writing the read L-Pages to respective physical addresses within the plurality of non-volatile memory devices;

updating the logical-to-physical address translation map for the valid entries to point to the respective physical addresses; and

generating new journal entries for the entries for which the logical-to-physical address translation map was updated.

11. The data storage device of claim 10, wherein selecting comprises weighing free space and program erase (PE) count in determining which S-Block to select.

12. The data storage device of claim 9, wherein the controller is further configured to garbage collect by at least:

selecting one of the plurality of S-Blocks to garbage collect;

reading the physical pages of the selected S-Block;

comparing L-Page numbers in the read physical pages of the selected S-Block to entries in the logical-to-physical translation map and designating entries that match as valid and entries that do not match as invalid;

writing the L-Pages corresponding to the valid entries to respective physical addresses within the plurality of non-volatile memory devices;

updating the logical-to-physical address translation map for the valid entries to point to the respective physical addresses; and

generating new journal entries for the entries for which the logical-to-physical address translation map was updated.

13. The data storage device of claim 1, wherein each journal number comprises a predetermined number of most significant bits of an address of a first physical page covered by the journal.

14. The data storage device of claim 1, wherein each of the plurality of journal entries comprises:

an L-Page number, and

a physical address location.

15. The data storage device of claim 1, wherein each of the plurality of journal entries comprises:

17

an L-Page number;
a physical address location of a physical page, and
an L-Page size.

16. The data storage device of claim 1, wherein each of the plurality of journal entries comprises:

a predetermined number of least significant bits of an address of a physical page that includes a start of an L-Page;
an address;
an L-Page size; and
an offset into the physical page.

17. The data storage device of claim 1, wherein the plurality of L-Pages are configured to be compressed and to vary in size, and wherein the plurality of journal numbers are configured to reference a greater number of L-Pages of smaller size or a lesser number of L-Pages of greater size.

18. The data storage device of claim 1, wherein the controller is further configured to read the plurality of journals upon startup in a predetermined sequential order and rebuild the logical-to-physical address translation map stored in volatile memory based upon the sequentially read plurality of journals.

19. The data storage device of claim 1, wherein the controller is further configured to build a journal map based on the plurality of journals.

20. The data storage device of claim 1, wherein the controller is further configured to:

read the plurality of journals upon startup in a predetermined sequential order;
build a map of the journals stored in the non-volatile memory devices based upon the sequentially read plurality of journals, and
store the built map of the journals in the volatile memory.

21. The data storage device of claim 1, wherein, among journal entries associated with a given L-Page, only a last-in-time updated journal entry associated with the given L-Page is valid.

22. The data storage device of claim 1, wherein the controller is further configured to maintain a system journal map of the plurality of journals in the volatile memory, each entry in the system journal map pointing to a location in the non-volatile memory devices where one of the plurality of journals is stored.

23. A method of controlling a data storage device comprising a volatile memory and a plurality of non-volatile memory devices, each of the plurality of non-volatile devices being configured to store a plurality of physical pages, each of the plurality of physical pages being stored at a predetermined physical location within the plurality of non-volatile devices, the method comprising:

storing data in a plurality of logical pages (L-Pages), each of the plurality of L-Pages being associated with an L-Page number that is configured to enable a controller to logically reference data stored in one or more of the physical pages;

maintaining a logical-to-physical address translation map in the volatile memory, the translation map being configured to enable determination of a physical location, within one or more of the physical pages, of the data stored in each L-Page, wherein an entry in the logical-to-physical address translation map includes an association between an L-page number and at least one of a length of the L-page or a size of error correcting code bits to be applied to the entry in the logical-to-physical address translation map;

maintaining a plurality of journals defining physical-to-logical correspondences in the plurality of non-volatile

18

memory devices, each of the plurality of journals being associated with a journal number, each journal covering a pre-determined range of physical pages and each comprising a plurality of journal entries, each entry being configured to associate one or more physical pages to each L-Page; and

reading the plurality of journals upon startup and rebuilding the logical-to-physical address translation map stored in volatile memory based upon the read entries in the plurality of journals.

24. The method of claim 23, further comprising creating a new entry in one of the plurality of journals upon an update to one of the plurality of L-Pages.

25. The method of claim 24, further comprising triggering write operations to the non-volatile memory devices to maintain a power-safe copy of translation data based on newly created journal entries instead of saving at least portions of the logical-to-physical address translation map, such that a write amplification is reduced.

26. The method of claim 24, wherein the new entry indicates a physical location, within a physical page, of a start of the updated L-Page.

27. The method of claim 24, further comprising updating a free space accounting by an amount corresponding to a length of the L-Page prior to being updated.

28. The method of claim 23, wherein storing comprises storing at least one of the plurality of L-Pages at a location that is unaligned with physical page boundaries.

29. The method of claim 23, wherein the physical pages are implemented as Error Correcting Code (ECC) pages (E-Pages) and wherein the plurality of devices comprises a plurality of flash memory blocks, each flash memory block comprising a plurality of flash memory pages (F-Pages), each of the F-Pages comprising a plurality of the E-Pages, each of the plurality of E-Pages being stored at a predetermined physical location within the plurality of devices.

30. The method of claim 23, further comprising updating the logical-to-physical address translation map with the physical location, within one or more physical pages, of the data referenced by the L-Page number of the updated L-Page.

31. The method of claim 23, wherein the plurality of devices comprises a plurality of super blocks (S-Blocks), each comprising one or more flash memory blocks per device and wherein each of the plurality of journal entries is configured to associated one or more of the physical pages of the S-Block to each L-Page.

32. The method of claim 31, further comprising:

selecting an S-Block to garbage collect;
comparing each entry in a journal for the selected S-Block to entries in the logical-to-physical address translation map and designating entries that match as valid and entries that do not match as invalid;

reading the L-Pages corresponding to the valid entries;
writing the read L-Pages to respective physical addresses within the plurality of non-volatile memory devices;
updating the logical-to-physical address translation map for the valid entries to point to the respective physical addresses; and
generating new journal entries for the entries for which the

logical-to-physical address translation map was updated.

33. The method of claim 32, wherein selecting comprises weighing free space and program erase (PE) count in determining which S-Block to select.

34. The method of claim 33, wherein the journal number comprises a predetermined number of most significant bits of an address of a first physical page covered by the journal.

19

35. The method of claim 31, further comprising selecting one of the plurality of S-Blocks to garbage collect; reading the physical pages of the selected S-Block; comparing L-Page numbers in the read physical pages of the selected S-Block to entries in the logical-to-physical address translation map and designating entries that match as valid and entries that do not match as invalid; writing the L-Pages corresponding to the valid entries to respective physical addresses within the plurality of non-volatile memory devices; updating the logical-to-physical address translation map for the valid entries to point to the respective physical addresses; and generating new journal entries for the entries for which the logical-to-physical address translation map was updated.

36. The method of claim 23, wherein each of the plurality of journal entries comprises:
an L-Page number, and
a physical address location.

37. The method of claim 23, wherein each of the plurality of journal entries comprises:
an L-Page number;
a physical address location of a physical page, and
an L-Page size.

38. The method of claim 23, wherein each of the plurality of journal entries comprises:
a predetermined number of least significant bits of an address of a physical page that includes a start of an L-Page;
an address;
an L-Page size; and
an offset into the physical page.

39. The method of claim 23, further comprising selectively compressing the plurality of L-Pages such that the plurality of L-Pages vary in size and wherein the plurality of journals are configured to reference a greater number of L-Pages of smaller size or a lesser number of L-Pages of greater size.

40. The method of claim 23, wherein reading the plurality of journals upon startup and rebuilding the logical-to-physical address translation map comprises reading the plurality of journals upon startup in a predetermined sequential order and rebuilding the logical-to-physical address translation map stored in volatile memory based upon the read plurality of journals.

41. The method of claim 23, wherein the controller is further configured to build a journal map based on the plurality of journals.

20

42. The method of claim 23, further comprising:
reading the plurality of journals upon startup in a predetermined sequential order;
building a map of the journals stored in the non-volatile memory devices based upon the sequentially read plurality of journals, and
storing the built map of the journals in the volatile memory.

43. The method of claim 23 wherein, among journal entries associated with a given L-Page, only a last-in-time updated journal entry associated with the given L-Page is valid.

44. The method of claim 23, further comprising maintaining a system journal map of the plurality of journals in the volatile memory, each entry in the system journal map pointing to a location in the non-volatile memory devices where one of the plurality of journals is stored.

45. A data storage device controller, comprising:
a processor configured to couple to a volatile memory and to a plurality of non-volatile memory devices, each of the plurality of non-volatile memory devices being configured to store a plurality of physical pages at a predetermined physical location within the plurality of non-volatile memory devices, the processor being further configured to program data to and read data from the plurality of non-volatile memory devices, the data being stored in a plurality of logical pages (L-Pages), each of the plurality of L-Pages being associated with an L-Page number that is configured to enable the processor to logically reference data stored in one or more of the physical pages, the volatile memory being configured to store a logical-to-physical address translation map configured to enable the processor to determine a physical location, within one or more physical pages, of the data stored in each L-Page,
wherein an entry in the logical-to-physical address translation map includes an association between an L-page number and at least one of a length of the L-page or a size of error correcting code bits to be applied to the entry in the logical-to-physical address translation map,
wherein the processor is configured to maintain, in the plurality of non-volatile memory devices, a plurality of journals defining physical-to-logical correspondences, each of the plurality of journals being associated with a journal number, each journal covering a pre-determined range of physical pages and comprising a plurality of journal entries, each entry being configured to associate one or more physical pages to each L-Page, wherein the processor is further configured to read the plurality of journals upon startup and rebuild the logical-to-physical address translation map stored in the volatile memory from the read plurality of journals.

* * * * *